

NVIDIA MLNX_OFED Documentation v24.07-0.6.1.0

Table of contents

Release Notes	7
General Support	10
Changes and New Features	23
Unsupported Functionalities/Features/NICs	26
Bug Fixes in This Version	26
Known Issues	29
User Manual	76
Introduction	76
Installation	87
Hardware and Software Requirements	88
Downloading the Drivers	88
Installing the Driver	90
Uninstalling the Driver	132
Updating Firmware After Installation	132
UEFI Secure Boot	136
Performance Tuning	139
Features Overview and Configuration	139
Ethernet Network	140
Ethernet Interface	141
Quality of Service (QoS)	144
Ethtool	160
Checksum Offload	164
Ignore Frame Check Sequence (FCS) Errors	165
RDMA over Converged Ethernet (RoCE)	165
Flow Control	180

	Explicit Congestion Notification (ECN)	189
	RSS Support	191
	Time-Stamping	193
	Flow Steering	203
	Wake-on-LAN (WoL)	208
	Hardware Accelerated 802.1ad VLAN (Q-in-Q Tunneling)	209
	VLAN Stripping in Linux Verbs	210
	Offloaded Traffic Sniffer	211
	Dump Configuration	211
	Local Loopback Disable	215
	Kernel Transport Layer Security (kTLS) Offloads	216
	IPsec Crypto Offload	218
	IPsec Full Offload	221
	MACsec Full Offload	228
lr	ifiniBand Network	232
	InfiniBand Interface	233
	NVIDIA SM	233
	QoS - Quality of Service	295
	IP over InfiniBand (IPoIB)	300
	Advanced Transport	313
	Optimized Memory Access	315
	NVIDIA PeerDirect	321
	CPU Overhead Distribution	322
	Out-of-Order (OOO) Data Placement	322
	IB Router	323
S	torage Protocols	324
	SRP - SCSI RDMA Protocol	325
	iSCSI Extensions for RDMA (iSER)	340
	Lustre	342
	NVME-oF - NVM Express over Fabrics	343

Virtualization	344
Single Root IO Virtualization (SR-IOV)	345
SR-IOV Live Migration	369
Enabling Paravirtualization	381
VXLAN Hardware Stateless Offloads	383
Q-in-Q Encapsulation per VF in Linux (VST)	384
802.1Q Double-Tagging	388
Scalable Functions	390
Resiliency	390
Reset Flow	390
Docker Containers	393
Docker Using SR-IOV	394
Kubernetes Using SR-IOV	394
Kubernetes with Shared HCA	395
HPC-X	395
Fast Driver Unload	395
OVS Offload Using ASAP ² Direct	395
Programming	479
InfiniBand Fabric Utilities	483
Troubleshooting	493
General Issues	494
Ethernet Related Issues	496
InfiniBand Related Issues	496
Installation Related Issues	497
Performance Related Issues	503
SR-IOV Related Issues	504
PXE (FlexBoot) Related Issues	504
RDMA Related Issues	505

Debugging Related Issues	505
OVS Offload Using ASAP2 Direct Related Issues	506
Common Abbreviations and Related Documents	506
Documentation History	510
Release Notes History	510
Changes and New Features History	510
Bug Fixes History	587
User Manual Revision History	658
Legal Notices and 3rd Party Licenses	660

Overview

i Note

MLNX_OFED has transitioned into DOCA-Host, and now available as DOCA-OFED (learn about DOCA-Host profiles <u>here</u>).

MLNX_OFED last standalone release is October 2024 Long Term Support (3 years). Starting January 2025 all new features will be included in DOCA-OFED only.

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:

- <u>Release Notes</u>
- User Manual

Software Download

 $\label{eq:please visit nvidia.com/en-us/networking \rightarrow \mbox{Products} \rightarrow \mbox{Software} \rightarrow \mbox{InfiniBand Drivers} \rightarrow \mbox{NVIDIA MLNX OFED}$

Document Revision History

For the list of changes made to the User Manual, refer to <u>User Manual Revision History</u>.

For the list of changes made to the Release Notes, refer to <u>Release Notes History</u>.

Release Notes

Release Notes Update History

Version	Date	Description
24.07-0.6.1.0		Initial release of this document version. This release introduces <u>Bug Fixes in This Version</u> .

j) 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:

- <u>General Support</u>
- Changes and New Features
- Bug Fixes in This Version
- <u>Known Issues</u>

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		 InfiniBand: SDR, FDR, EDR, HDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹
BlueField		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-7		 InfiniBand: EDR, HDR100, HDR, NDR200, NDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹, 200GbE ², 400GbE
ConnectX-6 Lx		• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ¹
ConnectX-6 Dx	mlx5	 Ethernet: 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹ , 200GbE ¹
ConnectX-6		 InfiniBand: SDR, FDR, EDR, HDR Ethernet: 10GbE, 25GbE, 40GbE, 50GbE ¹, 100GbE ¹ , 200GbE ¹
ConnectX- 5/ConnectX-5 Ex		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-4 Lx		• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE
ConnectX-4		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE ³, 100GbE

- 2. Speed that supports PAM4 mode only.
- 3. 56GbE is an NVIDIA proprietary link speed and can be achieved while connecting an NVIDIA adapter card to NVIDIA SX10XX switch series or when connecting an NVIDIA adapter card to another NVIDIA adapter card. ____

Embedded Components

Package	Revision	Licenses
clusterkit	1.14.462-1.2407061	BSD
dpcp	1.1.49-1.2407061	BSD-3-Clause
fwctl	24.07-OFED.24.07.0.5.7.1	GPLv2
hcoll	4.8.3228-1.2407061	Proprietary
ibarr	0.1.3-1.2407061	(GPL-2.0 WITH Linux-syscall-note) OR BSD-2-Clause
ibdump	6.0.0-1.2407061	BSD2+GPL2
ibsim	0.12-1.2407061	GPLv2 or BSD
ibutils2	2.1.1- 0.21800.MLNX20240804.ga435258 7.2407061	Mellanox Confidential and Proprietary
iser	24.07-OFED.24.07.0.5.7.1	GPLv2
isert	24.07-OFED.24.07.0.5.7.1	GPLv2
kernel-mft	4.29.0-131	Dual BSD/GPL
knem	1.1.4.90mlnx3-0FED.23.10.0.2.1.1	BSD and GPLv2
libvma	9.8.60-1	GPLv2 or BSD
libxlio	3.31.2-1	GPLv2 or BSD
mlnx-en	24.07-0.6.1.0.g51d9270	GPLv2
mlnx-ethtool	6.9-1.2407061	GPL
mlnx- iproute2	6.9.0-1.2407061	GPL

Package	Revision	Licenses
mlnx- nfsrdma	24.07-OFED.24.07.0.5.7.1	GPLv2
mlnx-nvme	24.07-0FED.24.07.0.5.7.1	GPLv2
mlnx- ofa_kernel	24.07-OFED.24.07.0.6.1.1	GPLv2
mlnx-tools	24.07-0.2407061	GPLv2 or BSD
mlx-steering- dump	1.0.0-0.2407061	GPLv2
mpitests	3.2.24-2ffc2d6.2407061	BSD
multiperf	3.0-3.0.2407061	BSD 3-Clause, GPL v2 or later
ofed-docs	24.07-OFED.24.07.0.6.1	GPL/BSD
ofed-scripts	24.07-0FED.24.07.0.6.1	GPL/BSD
openmpi	4.1.7a1-1.2407061	BSD
opensm	5.20.0.MLNX20240804.ef1f438a- 0.1.2407061	GPLv2 or BSD
openvswitch	2.17.8-1.2407061	ASL 2.0 and LGPLv2+ and SISSL
perftest	24.07.0-0.44.g57725f2.2407061	BSD 3-Clause, GPL v2 or later
rdma-core	2407mlnx52-1.2407061	GPLv2 or BSD
rshim	2.0.41-0.g5e402e1	GPLv2
sharp	3.8.0.MLNX20240804.aaa5caab- 1.2407061	Proprietary
sockperf	3.10-0.git5ebd327da983.2407061	BSD
srp	24.07-OFED.24.07.0.5.7.1	GPLv2
ucx	1.17.0-1.2407061	BSD
xpmem	2.7.3-1.2407061	GPLv2 and LGPLv2.1
xpmem-lib	2.7-0.2310055	LGPLv2.1

General Support

Supported Operating Systems

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
Alinux 3.2	x86 _64	5.10.134- 13.al8.x86_64	Prima ry	8	8	8	8	8	\bigotimes
Alma 8.5	x86 _64	4.18.0- 348.12.2.EL8_5 .X86_64	Com munit y	8	8	8	8	8	8
Anolis	AAr ch6 4	4.18.0- 348.2.1.AN8_4. AARCH64	Com munit y	8	8	8	8	8	8
OS 8.4	x86 _64	4.18.0- 305.AN8.X86_ 64	Com munit y	8	8	8	8	8	8
Anolis	AAr ch6 4	5.10.134+	Prima ry	8	8	8	8	8	8
OS 8.6	x86 _64	5.10.134+	Prima ry	8	8	8	8	8	8
BCLINU	AAr ch6 4	4.19.90- 2107.6.0.0098. oe1.bclinux.aar ch64	Prima ry	8	8	8	•	8	8
X21.10 SP2	x86 _64	4.19.90- 2107.6.0.0100. oe1.bclinux.x8 6_64	Prima ry	8	8	•	0	8	8
BCLinux 22.10	aarc h64	5.10.0- 153.24.0.100.6. oe2203sp2.bcli nux.aarch64	Prima ry	8	8	8	⊘	8	8
	x86	5.10.0- 153.24.0.100.6.	Prima ry	8	8	•	0	8	8

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
		oe2203sp2.bcli nux.x86_64							
CentOS	AAr ch6 4	4.18.0- 535.el8.aarch6 4	Com munit y	8	8	8	8	8	8
Stream v8	x86 _64	4.18.0- 535.el8.x86_64	Com munit y	8	8	8	8	8	8
CentOS	AAr ch6 4	5.14.0- 407.el9.x86_64	Com munit y	8	8	8	8	8	8
Stream v9	x86 _64	5.14.0- 407.el9.aarch6 4	Com munit y	8	8	8	8	8	8
CTYUN	AAr ch6 4	4.19.90- 2102.2.0.0062. ctl2.aarch64	Prima ry	8	8	8	8	8	8
OS2.0	x86 _64	4.19.90- 2102.2.0.0062. ctl2.x86_64	Prima ry	8	8	8	8	8	8
CTYUN	AAr ch6 4	5.10.0- 136.12.0.86.ctl 3.aarch64	Prima ry	8	8	8	8	8	8
0S23.0 1	x86 _64	5.10.0- 136.12.0.86.ctl 3.x86_64	Prima ry	8	8	8	8	8	8
Debian 1 0.8	AAr ch6 4	4.19.0-14- arm64	Prima ry	8	8	8	8	8	8
Debian 1 0.9	x86 _64	4.19.0-14- amd64	Prima ry	8	8	8	8	8	⊗

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
	x86 _64	4.19.0-16- amd64	Prima ry	•	8	8		8	8
Debian 1	AAr ch6 4	4.19.0-21- arm64	Prima ry	8	8	8	⊗	8	8
0.13	x86 _64	4.19.0-21- amd64	Prima ry	8	8	8	8	8	8
Debian 1	AAr ch6 4	5.10.0-13- arm64	Prima ry	Ø	8	8	0	8	8
1.3	x86 _64	5.10.0-13- amd64	Prima ry	0	⊗	8	0	8	8
Debian 1	AAr ch6 4	6.1.0-10- arm64	Prima ry	0	8	8	0	8	8
2.1	x86 _64	6.1.0-10- amd64	Prima ry	0	8	8	•	8	
Debian 1	aarc h64	6.1.0-18- arm64	Prima ry	•	8	8	•	8	8
2.5	x86	6.1.0-18- amd64	Prima ry	0	8	8	•	8	8
EulerOS	AAr ch6 4	4.19.90- vhulk2006.2.0. h171.eulerosv2 r9.aarch64	Com munit y	8	8	8	8	8	8
2.0sp9	x86 _64	4.18.0- 147.5.1.0.h269. eulerosv2r9.x8 6_64	Com munit y	8	8	8	8	8	8
EulerOS 2.0sp10	AAr ch6 4	4.19.90- vhulk2110.1.0.	Com munit y	8	8	8	8	8	8

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
		h860.eulerosv2 r10.aarch64							
	x86 _64	4.18.0- 147.5.2.4.h694. eulerosv2r10.x 86_64	Com munit y	8	8	8	8	8	8
EulerOS	AAr ch6 4	5.10.0- 60.18.0.50.h32 3.eulerosv2r11. aarch64	Prima ry	8	8	8	⊗	8	⊗
2.0sp11	x86 _64	5.10.0- 60.18.0.50.h32 3.eulerosv2r11. x86_64	Prima ry	8	8	⊗	⊗	8	⊗
EulerOS	AAr ch6 4	5.10.0- 136.12.0.86.h1 032.eulerosv2r 12.aarch64	Prima ry	8	8	8	8	8	8
2.0sp12	x86 _64	5.10.0- 136.12.0.86.h1 032.eulerosv2r 12.x86_64	Prima ry	8	8	8	8	8	8
KYLIN 1	AAr ch6 4	4.19.90- 24.4.v2101.ky1 0.aarch64	Prima ry	8	8	8	8	8	8
OSP2	x86 _64	4.19.90- 24.4.v2101.ky1 0.x86_64	Prima ry	Ø	8	8	8	8	8
KYLIN1 OSP3	AAr ch6 4	4.19.90- 52.22.v2207.ky 10.aarch64	Prima ry	Ø	8	8	8	8	8
	x86 _64	4.19.90- 52.22.v2207.ky	Prima ry	•	8	8	8	8	8

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
		10.x86_64							
Mariner 2.0	x86 _64	5.15.148.2- 2.cm2	Prima ry	⊗	8	8	8	\bigotimes	8
Oracle Linux 7.9	x86 _64	5.4.17- 2011.6.2.el7ue k.x86_64	Prima ry	⊗	8	⊗	⊗	8	8
Oracle Linux 8.4	x86 _64	5.4.17- 2102.201.3.el8 uek.x86_64	Prima ry	8	8	⊗	⊗	8	8
Oracle Linux 8.6	x86 _64	5.4.17- 2136.307.3.1.el 8uek.x86_64	Prima ry	8	8	8	8	8	8
Oracle Linux 8.7	x86 _64	5.15.0- 3.60.5.1.el8uek .x86_64	Prima ry	8	8	⊗	⊗	8	8
Oracle Linux 8.8	x86 _64	5.15.0- 101.103.2.1.el8 uek.x86_64	Prima ry	8	8	8	8	8	8
Oracle Linux 9.0	x86 _64	5.15.0- 0.30.19.el9uek. x86_64	Prima ry	8	8	⊗	⊗	8	8
Oracle Linux 9.1	x86 _64	5.15.0- 3.60.5.1.el9uek .x86_64	Prima ry	8	8	⊗	⊗	8	⊗
Oracle Linux 9.2	x86 _64	5.15.0- 101.103.2.1.el9 uek.x86_64	Prima ry	8	8	8	8	8	8
OpenSU SE 15.3	AAr ch6 4	-	Com munit y	8	8	8	8	8	8
	x86 _64	5.3.18- 150300.59.43-	Com munit	8	8	8	8	8	8

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
		DEFAULT	У						
OPENE ULER20	AAr ch6 4	4.19.90- 2012.4.0.0053. OE1.AARCH64	Com munit y	8	8	⊗	8	⊗	⊗
.03SP1	x86 _64	4.19.90- 2110.8.0.0119. OE1.X86_64	Com munit y	8	8	8	8	8	⊗
	AAr ch6 4	4.19.90- 2112.8.0.0131. oe1.aarch64	Prima ry	8	8	8	8	8	8
ULER20 .03SP3	x86 _64	4.19.90- 2112.8.0.0131. oe1.x86_64	Prima ry	8	8	8	8	8	8
	AAr ch6 4	5.10.0- 60.18.0.50.oe2 203.aarch64	Prima ry	8	8	8	8	8	8
ULER22 .03	x86 _64	5.10.0- 60.18.0.50.oe2 203.x86_64	Prima ry	8	8	8	8	8	8
openEul er 22.03 SP1	x86	5.10.0- 136.12.0.86.oe 2203sp1.x86_ 64	Prima ry	8	8	8	8	8	8
Photon OS 3.0	x86 _64	4.19.225-3.ph3	Com munit y	8	8	8	8	8	⊗
RHEL/C entOS8.	AAr ch6 4	4.18.0- 80.el8.aarch64	Prima ry	0	0	8	8	0	12.5
0	x86 _64	4.18.0- 80.el8.x86_64	Prima ry	•	O	8	8	•	12.5

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
RHEL/C entOS8.	AAr ch6 4	4.18.0- 147.el8.aarch6 4	Prima ry	•	•	8	8	•	12.5
1	x86 _64	4.18.0- 147.el8.x86_64	Prima ry	0	0	8	8	0	12.5
RHEL/C entOS8.	AAr ch6 4	4.18.0- 193.el8.aarch6 4	Prima ry	•	Ø	8	0	Ø	12.5
2	x86 _64	4.18.0- 193.el8.x86_64	Prima ry	0	0	0	0	0	12.5
RHEL/C entOS8.	AAr ch6 4	4.18.0- 240.el8.aarch6 4	Prima ry	•	0	8	8	•	12.5
3	x86 _64	4.18.0- 240.el8.x86_64	Prima ry	0	◙	8	8	0	12.5
RHEL/C entOS8.	AAr ch6 4	4.18.0- 305.el8.aarch6 4	Prima ry	•	•	8	•	•	12.5
4	x86 _64	4.18.0- 305.el8.x86_64	Prima ry	0	•	•	•	0	12.5
RHEL/C entOS/	AAr ch6 4	4.18.0- 348.el8.aarch6 4	Prima ry	◙	0	8	⊘	◙	12.5
Rocky8. 5	x86 _64	4.18.0- 348.el8.x86_64	Prima ry	•	⊘		⊘	•	12.5
RHEL/R	AAr ch6 4	AArch644.18.0 - 372.41.1.el8_6. aarch64	Prima ry	0	Ø	8	•	0	12.5
ocky8.6	x86 _64	4.18.0- 372.41.1.el8_6. x86_64	Prima ry	0	0	0		0	12.5

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
RHEL/R	AAr ch6 4	4.18.0- 425.14.1.el8_7. aarch64	Prima ry	•	8	8	•	v	12.5
ocky8.7	x86 _64	4.18.0- 425.14.1.el8_7. x86_64	Prima ry	•	8	Ø	•	v	12.5
RHEL/R	AAr ch6 4	4.18.0- 477.10.1.el8_8. aarch64	Prima ry	•	8	8	0	v	12.5
ocky8.8	x86 _64	4.18.0- 477.10.1.el8_8. x86_64	Prima ry	•	8	Ø	0	•	12.5
RHEL/R	AAr ch6 4	4.18.0- 513.5.1.el8_9.a arch64	Prima ry	•	8	8	0	v	12.5
ocky8.9	x86 _64	4.18.0- 513.5.1.el8_9.x 86_64	Prima ry	•	8	Ø	0	•	12.5
RHEL/R	AAr ch6 4	5.14.0- 70.46.1.el9_0.a arch64	Prima ry	•	8	8	0	•	12.5
ocky9.0	x86 _64	5.14.0- 70.46.1.el9_0.x 86_64	Prima ry	•	8	Ø	0	v	12.5
RHEL/R	AAr ch6 4	5.14.0- 162.19.1.el9_1. aarch64	Prima ry	Ø	8	8	0	Ø	12.5
ocky9.1	x86 _64	5.14.0- 162.19.1.el9_1. x86_64	Prima ry	0	8	Ø	0	•	12.5
RHEL/R ocky9.2	AAr ch6	5.14.0- 284.11.1.el9_2.	Prima ry	0	8	8	•	0	12.5

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
	4	aarch64							
	x86 _64	5.14.0- 284.11.1.el9_2. x86_64	Prima ry	•	8	Ø	0	•	12.5
RHEL/R	AAr ch6 4	5.14.0- 362.8.1.el9_3.a arch64	Prima ry	◙	8	•	⊘	◙	12.5
ocky9.3	x86 _64	5.14.0- 362.8.1.el9_3.x 86_64	Prima ry	v	8	Ø	0	•	12.5
RHEL/R	AAr ch6 4	5.14.0- 427.13.1.el9_4. aarch64	Prima ry	•	8	Ø	0	•	12.5
ocky9.4	x86 _64	5.14.0- 427.13.1.el9_4. x86_64	Prima ry	Ø	8	ø	0	•	12.5
RHEL/R	aarc h64	4.18.0- 553.el8_10.aar ch64	Prima ry	•	8	8	0	•	12.5
ocky 8.10	x86	4.18.0- 553.el8_10.x86 _64	Prima ry	•	8	•	•	•	12.5
SLES15	AAr ch6 4	5.3.18-22- default	Prima ry	◙	•	8	⊘	8	⊗
SP2	x86 _64	5.3.18-22- default	Prima ry	0	•	•	0	8	8
SLES15 SP3	AAr ch6 4	5.3.18-57- default	Prima ry	0	8	8	•	8	⊗
573	x86 _64	5.3.18-57- default	Prima ry	•	8	0	0	⊗	⊗

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
SLES15	AAr ch6 4	5.14.21- 150400.22- default	Prima ry	⊘	8	8	•	8	8
SP4	x86 _64	5.14.21- 150400.22- default	Prima ry	•	8	Ø	•	8	8
SLES15	AAr ch6 4	5.14.21- 150500.53- default	Prima ry	0	8	8	0	8	8
SP5	x86 _64	5.14.21- 150500.53- default	Prima ry	•	8	Ø	0	8	8
SLES 15 SP6	x86	6.4.0- 150600.21- default	Prima ry	Ø	8	Ø	0	8	8
tencent	AAr ch6 4	5.4.119- 19.0009.39	Prima ry	8	8	8	0	8	8
os3.3	x86 _64	5.4.119- 19.0009.39	Prima ry	8	8	8	0	8	8
Ubuntu	AAr ch6 4	5.4.0-26- generic	Prima ry	0	0	8	0	0	12.5
20.04	x86 _64	5.4.0-26- generic	Prima ry	0	•	0	◙	0	12.5
Ubuntu	AAr ch6 4	5.15.0-25- generic	Prima ry	0	Ø	Ø	0	0	12.5
22.04	x86 _64	5.15.0-25- generic	Prima ry	0	⊘	Ø	◙	0	12.5
Ubuntu 24.04	AAr ch6	6.6.0-14- generic	Prima ry		0	•	•	•	12.5

Operati ng System	Arc hite ctur e	Default Kernel Version (Primary)/ Tested with Kernel Version (Community)	OS Suppo rt Model	ASAP2 OVS- Kernel SR-IOV	ASAP2 OVS- DPDK SR-IOV	NFS over RDM A	NV M e	GPUDir ect Storage (GDS)	UCX- CUDA Versio n
	4								
	x86 _64	6.6.0-14- generic	Prima ry	•	0		0	0	12.5
UOS20.	AAr ch6 4	5.10.0- 46.uel20.aarch 64	Prima ry	⊗	8	8	8	8	⊗
1060	x86 _64	5.10.0- 46.uel20.x86_6 4	Prima ry	8	8	8	8	8	8
UOS20.	AAr ch6 4	5.10.0- 46.uelc20.aarc h64	Prima ry	8	8	8	8	8	8
1060a	x86 _64	5.10.0- 46.uelc20.x86_ 64	Prima ry	8	8	8	8	8	8
Citrix XenSer ver Host8.2	x86 _64	4.19.0+1	Prima ry	8	8	8	8	8	8
Linux	aarc h64	6.10	Prima ry	•	8	8	0	8	8
Kernel 6.10	x86	0.10	Prima ry	•	8	0	0	8	⊗

(i) 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
24.07-0.6.1.0 July 2024	24.04-0.7.0.0 - DOCA-OFED Profile	GA	April 2024
	23.10-3.2.2.0 - MLNX_OFED	LTS-Update	June 2023
	5.8-5.1.1.2	LTS-Update	June 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
24.07-0.6.1.0 July 2024	24.04-0.7.0.0 - DOCA-OFED Profile	GA	April 2024
	5.8-5.1.1.2	LTS-Update	June 2023

Supported NIC Firmware Versions

(i) 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	24.42.1000
ConnectX-7	28.42.1000
ConnectX-6 Lx	26.42.1000
ConnectX-6 Dx	22.42.1000
ConnectX-6	20.42.1000
ConnectX-5/ConnectX-5 Ex	16.35.4030
BlueField	18.33.1048
ConnectX-4 Lx	14.32.1010
ConnectX-4	12.28.2006

For the official firmware versions, please see <u>https://www.nvidia.com/en-us/networking/</u> \rightarrow Support \rightarrow Support \rightarrow Firmware Download.

Changes and New Features

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

Supported Cards	Description
All HCAs	Supported in the following adapter cards <u>unless specifically</u> <u>stated otherwise:</u> 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 <u>unless specifically</u> <u>stated otherwise:</u> ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2
ConnectX-6 and above	Supported in the following adapter cards <u>unless specifically</u> <u>stated otherwise:</u> ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX- 7/BlueField-2
ConnectX-5 and above	Supported in the following adapter cards <u>unless specifically</u> <u>stated otherwise:</u> ConnectX-5/ConnectX-6/ConnectX-6 Dx/ConnectX-6 Lx/ConnectX-7/BlueField-2
ConnectX-4 and above	Supported in the following adapter cards <u>unless specifically</u> <u>stated otherwise:</u> 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 <u>Release Notes Change Log History</u> section.

Feature/Change	Description						
	24.07-0.6.1.0						
SFs On-Demand EQs	Changed the SF initiation/boot flow to avoid EQ creations. Now EQ will be created only when required by the user/system QPs.						
fwctl Subsystem	[Beta] fwctl is a new subsystem intended to bring some common rules and order to the growing pattern of exposing a secure firmware interface directly to the userspace. fwctl focuses on debugging, configuration and provisioning of the device. As fwctl is not installed by default, you can install it using one of the following ways:						
	 Using the MLNX_OFED script by adding "fwctl" to the installation command. Using the package manager, after installing MLNX_OFED or DOCA-Host: search for the fwctl package: 						

Feature/Change	Description	
	 yum/apt/zypper search fwctl install the result from the last command. 	
	yum/apt/zypper install <name fwctl<br="" of="">package from last command></name>	
Exposing req_transport_ret ries_exceeded Counter	Exposes the req_transport_retries_exceeded counter (per QP counter). It can be seen using the RDMA statistics, and shows the number of times the requester detected transport retries exceed error.	
Socket-Direct multi-PF netdev	Creates a single netdev abstraction for multiple Socket-Direct PFs of the same port that run on multiple NUMA nodes. This single netdev handles traffic locally on each of the NUMA nodes without upper layer complexities. This improves the applications' locality and performance in a multi-NUMA system.	
Synchronous (Scalable Function) SF Activation for Local SFs	Added devlink-nested option to show relationships between devlink instances through devlink dev show/devlink port show commands. The devlink port show is used to check when SFs are activated, and shows the relationship between SF auxiliary device devlink instance and the eSwitch representor devlink port. devlink Scalable Function (SF) activate command is received before enabling a partial parallelization of the SFs activation, and reduces time.	
General	Bug fixes	

Customer Affecting Changes

Customer Affecting Change Description	
Tools: mstflint	As of MLNX_OFED v24.07, mstflint will no longer be installed, upgraded, or removed as part of the package installation. To download it, go to <u>https://github.com/Mellanox/mstflint</u> .

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 (libibervs, libmlx5, etc.) to the Upstream version rdma-core.

For the list of MLNX_OFED verbs APIs that have been migrated, refer to <u>Migration to</u> <u>RDMA-Core document</u>.

Unsupported Functionalities/Features/NICs

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

- Relational Database Service (RDS)
- mthca InfiniBand driver
- Ethernet IPoIB (eIPoIB)
- InfiniBand Connected transport service
- IPSec, MACsec, and TLS acceleration over bond for crypto offload

Bug Fixes in This Version

Below are the bugs fixed in this version. For a list of fixes previous version, see <u>Bug Fixes</u> <u>History</u>.

Internal Reference Number	Description
3916656	Description: If multicast join request is received during the Heavy Sweep, the InfiniBand Subnet Manager 5.19 may hang.
	Keywords: Multicast, Subnet Manager

Internal Reference Number	Description
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed an issue where executing cma_roce_mode -d [device] would inadvertently delete contents of /sys/kernel/config/rdma_cm. This occurred due to the removal of the rdma_cm/ device directory upon script exit, affecting system configuration.
3878076	Keywords: cma_roce_mode configuration files deletion
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed an issue that resulted in corrupt SMP MAD requests list when the sent list was accessed while the unregistered flow was running.
3893526 /	Keywords: SMP MAD requests
3869558	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed the packet inspection parsing to avoid data corruption when GRE offload was turned on by parsing the outer header as UDP and not as TCP.
3896666	Keywords: UDP, TCP, GRE offload
	Discovered in Release: 24.01-0.3.3.1
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed an issue that prevented netdev queue value from being updated in mqprio param when switchdev mode was enabled and the netdev queue number was reset to 1.
3974743	Keywords: netdev
	Discovered in Release: 23.10-0.5.5.0
	Fixed in Release: 24.07-0.6.1.0
3974743	Description: Fixed the error flow in case of performing CT entry update which resulted in failure as it exceeded the firmware limit.
	Keywords: Offload
	Discovered in Release: 24.04-0.6.6.0

Internal Reference Number	Description
	Fixed in Release: 24.07-0.6.1.0
	Description: Extended the number of rules to be offloaded to all ports from 8 to 32. This allows greater flexibility when offloading temporary rules like VLAN traffic and ease the offload sharing between the ports.
3957895	Keywords: Offload rules, VLAN traffic
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed a stack overrun warning by reducing the size of the local on-stack array used for optimization by 192 bytes.
3936207	Keywords: Kernel stack
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed an overflow issue where RDMA applications could not find the device when creating new RDMA devices.
3855582	Keywords: RDMA CM
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed the setting of ATS for DMABUF MRs that caused some MRs to miss the ATS enablement. Lack of ATS enablement on DMABUF MRs results in slower performance when using these MRs.
3932944	Keywords: ATS
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0
	Description: Fixed an issue related to the driver's internal MR cache cleanup process that resulted in high memory consumption on the host.
3947203	Keywords: MR cache
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.07-0.6.1.0

Known Issues

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

Internal Ref. Number	Issue
	Description: DOCA-HOST include drivers for both Ethernet and InfiniBand protocols. By installing the kernel part of DOCA-HOST (as with MLNX_OFED before it) the InfiniBand drivers, such as ib_core and its dependencies that you may have previously installed on your host server will be replaced with the DOCA-HOST drivers.
3990230	Keywords: DOCA-HOST, kernel, inbox drivers
	Workaround: To maintain the non-DOCA drivers, please use the relevant Inbox drivers <u>https://docs.nvidia.com/networking/software/adapter-</u> software/index.html#linux-inbox-drivers-upstream-releases
	Discovered in Release: 24.07-0.6.1.0
	Description: An mlx5_core probe error may occur on RHEL 9.4 after unloading the mlx5_core module.
4040107	Keywords: mlx5_core
4040187	Workaround: Upgrade to kernel 5.14.0-427.20.1.el9_4 or newer to resolve the kernel panics.
	Discovered in Release: 24.07-0.6.1.0
	Description: sFlow is not supported in MLNX_OFED v24.07-x.x.x.x.
4020072	Keywords: sFlow
4030072	Workaround: N/A
	Discovered in Release: 24.07-0.6.1.0
4022803	Description: fwctl subsystem is supported on default kernels only. "add-kernel- support" will not build and install fwctl.
	Keywords: fwctl, kernel
	Workaround: N/A
	Discovered in Release: 24.07-0.6.1.0

Internal Ref. Number	Issue
	Description: When using QEMU versions older than 8.2, interruption loss are experienced resulting in firmware commands timeouts and undefined behavior in the driver.
	Keywords: QEMU
	Workaround: To avoid this issue, make sure:
4001184	 the running kernel includes the following change, which was accepted in 6.4-rc3: [Kernel Patch] (https://lore.kernel.org/lkml/28521e1b0b091849952b0ecb8c118729fc8 cdc4f.1683740667.git.reinette.chatre@intel.com/T/) QEMU version is 8.2 or newer: [Qemu Patch] (https://lore.kernel.org/qemu-devel/20231009064900.1465361-5- clg@redhat.com/)
	Discovered in Release: 24.07-0.6.1.0
	Description: In Debian 12, using dhcpcd instead of dhclient to configure the network interface (using Networkmanager) will result in wrong network interface configuration.
3856101	Keywords: dhcpcd, dhclient, Debian 12, Networkmanager
	Workaround: Use dhclient to configure the network interface.
	Discovered in Release: 24.04-0.6.6.0
	Description: Driver might try to access privileged registers resulting in an error with syndrome.
3964215	Keywords: Unbind and bind the function or restart the driver.
	Workaround: N/A
	Discovered in Release: 24.04-0.6.6.0
3640907	Description: 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.
	Keywords: PCIe Gen5, IOMMU, Call Trace
	Workaround: To resolve the issue either:
	 Add kernel parameter cmdline "iommu=pt"

Internal Ref. Number	Issue
	or • Use a kernel that includes this upstream patchset: <u>https://patchwork.kernel.org/project/linux-</u> <u>mediatek/cover/20190908165642.22253-1-murphyt7@tcd.ie/</u>
	Discovered in Release: 24.04-0.6.6.0
	Description: 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.
3004304	Keywords: NVMe, Call Trace, num_p2p_queues
	Workaround: N/A
	Discovered in Release: 24.01-0.3.3.1
	Description: 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 .
3735400	Keywords: NVMF connect, PI, Adaptive Routing
5155166	Workaround: Disable the AR at the opensm, or, alternatively, disable the PI at the nvme_rdma with a new module parameter.
	Discovered in Release: 24.01-0.3.3.1
	Description: 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.
3774149	Keywords: Race condition, RDMA_WRITE, shared memory write
	Workaround: Set UCX_RNDV_SCHEME=get_zcopy to force using RDMA_READ protocol.
	Discovered in Release: 24.01-0.3.3.1
3565433	Description: 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.
	Keywords: DCI, SGEs
	Workaround: N/A

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

Internal Ref. Number	Issue
	Description: 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.
3678715	Keywords: openIbd service, nvme_rdma module
	Workaround: Manually unload the nvme_rdma module before performing the driver restart. This can be achieved using the modprobe -r nvme_rdma command.
	Discovered in Release: 23.10-1.1.9.0
3676223	Description: When using kernel version 4.12 or above, it is advised to run echo 0 > /sys/bus/pci/devices/0000\:08\:00.0/sriov_drivers_autoprobe to avoid VF probing

Internal Ref. Number	Issue
	Keywords: VF probing
	Workaround: N/A
	Discovered in Release: 23.10-1.1.9.0
2002050	Description: 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.
3682658	Keywords: RDMA-CM user application, AF_IB, private data
	Workaround: Do not use AF_IB for application's private data.
	Discovered in Release: 23.10-0.5.5.0
3640082	 Description: 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."
	Keywords: Maximal VF number
	Workaround: N/A
	Discovered in Release: 23.10-0.5.5.0
	 Description: When offloading IPsec policy rules while in legacy mode there are two options: 1. Software steering - The software stack will handle the task, and no device offload will take place.
3653417	 Changing the steering mode to firmware steering will return unsupported.
	Keywords: IPsec, legacy mode
	Workaround: Perform a devlink reload after changing the steering mode.
	Discovered in Release: 23.10-0.5.5.0
3612274	Description: 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.
	Keywords: IPsec offload, TC offload

Internal Ref. Number	Issue
	Workaround: N/A
	Discovered in Release: 23.10-0.5.5.0
	Description: 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.
3596126	Keywords: OVS mirroring, Connectx-5
	Workaround: N/A
	Discovered in Release: 23.10-0.5.5.0
2522.422	Description: 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 https://www.suse.com/support/kb/doc/?id=000021137.
3538463	Keywords: Kernel ABI, Sles15SP4, driver start
	Workaround: Upgrade to a kernel version newer than 5.14.21-150400.24.63.1 (May 2023).
	Discovered in Release: 23.10-0.5.5.0
	Description: When running over REHL7.6 with excessive RDMA/RoCE workload, kernel warnings may be triggered.
3637252	Keywords: REHL7.6, RDMA, RoCE
	Workaround: N/A
	Discovered in Release: 23.10-0.5.5.0

Internal Ref. Number	Issue
3046655	Description: A package manager upgrade with zypper (on an SLES system) may prompt a question about vendor change from "Mellanox Technologies" to "OpenFabrics".
	Keywords: Installation, SLES
	Workaround: Either accept the prompted change, or add the /etc/zypp/vendors.d/mlnx_ofed file with the following content: [main] vendors = Mellanox,OpenFabrics

Internal Ref. Number	Issue
	Discovered in Release: 23.07-0.5.0.0
3392477	Description: The ConnectX-7 firmware embedded in this MLNX_OFED version cannot be burnt using the MLNX_OFED installer script.
	Keywords: ConnectX-7, MLNX_OFED installer script
	Workaround: Please download and install the dedicated firmware from the web <u>https://network.nvidia.com/support/firmware/connectx7ib/</u>
	Discovered in Release: 23.07-0.5.0.0
3532756	Description: The kernel may crash when restarting the driver while IP sec rules are configured.
	Keywords: IP sec
	Workaround: Flush the IP sec configuration before reloading the driver: ip xfrm state flush ip xfrm policy flush
	Discovered in Release: 23.07-0.5.0.0
3472979 3413938	Description: When a large number of virtual functions are present, the output of the "ip link show" command may be truncated.
	Keywords: virtual functions, ip link show
	Workaround: N/A
	Discovered in Release: 23.07-0.5.0.0
	Description: 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.
	Keywords: Hang; mlnx-sf
	Workaround: N/A
	Discovered in Release: 23.07-0.5.0.0
3461572	Description: 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.
	Keywords: Multiport Eswitch, compat sysfs, mpesw LAG

Internal Ref. Number	Issue
	Workaround: N/A
	Discovered in Release: 23.07-0.5.0.0
	Description: Simultaneously adding or removing TC rules while operating on kernel version 6.3 could potentially result in stability issues.
2 4 6 4 2 2 7	Keywords: ASAP, rules, TC
3464337	Workaround: Make sure the following fix is part of the kernel: <u>https://lore.kernel.org/netdev/20230504181616.2834983-3-</u> <u>vladbu@nvidia.com/T/</u>
	Discovered in Release: 23.07-0.5.0.0
	Description: 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 test, the actual number of offloaded connections may be lower than expected.
	Keywords: ASAP, CT offload
3469484	Workaround: Make sure to have the following offending commit in the tree: net/sched: act_ct: offload UDP NEW connections Make sure to to have <u>https://www.spinics.net/lists/stable-</u> <u>commits/msg303536.html</u> in the kernel tree to fix this issue.
	Discovered in Release: 23.07-0.5.0.0
3473331	Description: When performing a CT offload test, the actual number of offloaded connections may be lower than expected.
	Keywords: ASAP, CT offload
	Workaround: 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: <u>https://www.spinics.net/lists/stable-commits/msg303536.html</u> in the kernel tree to fix this issue.
	Discovered in Release: 23.07-0.5.0.0

Internal Ref. Number	Issue
	Description: Configuring PFC in parallel to buffer size and prio2buffer commands may lead to misalignment between firmware and software in regards to receiving buffer ownership.
220710	Keywords: NetDev, PFC, Buffer Size, prio2buffer
3360710	Workaround: First, configure PFC on all ports, and then perform other needed QoS (i.e., buffer_size or prio2buffer) configurations accordingly.
	Discovered in Release: 23.04-0.5.3.3
	Description: OpenSM may not be started automatically if chkconfig was not installed before OpenSM is installed. Note, however, that 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).
3413879	Keywords: Installation, OpenSM, chkconfig
	Workaround: Install chkconfig before installing MLNX_OFED. If installing it fails, make sure /etc/init.d does not exist at the time of installing it.
	Discovered in Release: 23.04-0.5.3.3
	Description: On SLES 15.4, installing MLNX_OFED using a package repository (with zypper) may trigger an error message about missing dependency for ' <u>librte_eal.so</u> .20.0()(64bit)'. This is because the inbox package libdpdk-20_0 is being uninstalled as it is incompatible with the MLNX_OFED rdma-core packages.
3424596	Keywords: Installation, SLES 15.4
	Workaround: Uninstall the relevant packages: 'zypper uninstall libdpdk-20_0' before installing MLNX_OFED. This will also remove the inbox openvswitch package.
	Discovered in Release: 23.04-0.5.3.3
3433416	Description: 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.
	Keywords: Installation, CUDA, yum

Internal Ref. Number	Issue
	Workaround: 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.
	Discovered in Release: 23.04-0.5.3.3
	Description: mlx-steering-dump is not supported on systems in which Python3 is not the default.
3420831	Keywords: mlx-steering-dump, Python3
	Workaround: N/A
	Discovered in Release: 23.04-0.5.3.3
	Description: 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").
	Keywords: Persistant Interface Names
3351989	Workaround: Add thecopy-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.
	Discovered in Release: 23.04-0.5.3.3

Internal Ref. Number	Issue
	Description: 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.
3324094	Keywords: NetDev, MTU, UDP Stream
	Workaround: N/A
	Discovered in Release: 5.9-0.5.6.0
3313137	Description: 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.

Internal Ref. Number	Issue
	This is the standard widely seen behavior in Linux drivers today.
	Keywords: Core, SR-IOV, VF, PF
	Workaround: N/A
	Discovered in Release: 5.9-0.5.6.0
	Description: When the system is overloaded, there is a possibility that one hour will pass between the creation of DevLink port and it usage/assignment, due to some locking. This will trigger a trace starting with: "Type was not set for devlink port."
3320947	Keywords: Core, DevLink, System Overload
	Workaround: N/A
	Discovered in Release: 5.9-0.5.6.0
	Description: 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.
3046222	Keywords: Installation, Ubuntu22
	Workaround: Usewith-openvswitch flag along with the installation command.
	Discovered in Release: 5.9-0.5.6.0
	Description: Devlink reload while deleting namespace may causes a deadlock on kernels older than Linux-6.0.
3262725	Keywords: Devlink, Namespace
	Workaround: N/A
	Discovered in Release: 5.9-0.5.6.0
3253255	 Description: 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).

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

Internal Ref. Number	Issue
	Description: On EulerOS 2.0SP11, installation with the yum method may fail with an error that mlnx-iproute2 is missing a dependency on <u>libdb-5.3.so()(64bit)</u> .
3215514	Keywords: Installation, EulerOS 2.0SP11, yum
5215514	Workaround: Install in advance the mlnx-iproute2 package with rpm and with thenodeps option. For example: rpm -Uvnodeps RPMS/mlnx-iproute2-5.19.0-1.58101.x86_64.rpm
	Discovered in Release: 5.8- 1.0.1.1
	Description: 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.
3191223	Keywords: ASAP ² , eSwitch, TC Rules
	Workaround: Set eSwitch mode to legacy before enabling SR-IOV or reload mlx5_core to change eSwitch mode to legacy.
	Discovered in Release: 5.8- 1.0.1.1
	Description: ping -6 -i <interface name=""> is broken in v5.18.</interface>
	Keywords: NetDev, -i flag
3199628	Workaround: In all operating systems that are running Kernel 5.18 and below, remove the -i flag.
	Discovered in Release: 5.8- 1.0.1.1
3002932	Description: Jumbo MTU must be set on all uplinks (i.e., uplinks of *_sf and *_sf_r) at all times.
	Keywords: NetDev, MTU, Uplink
	Workaround: Configure jumbo MTU (9216) on all uplink-related interfaces.

Internal Ref. Number	Issue
	Discovered in Release: 5.8- 1.0.1.1
	Description: The yum install method might be broken on installer regenerated withadd-kernel-support-build-only.
2120050	Keywords: Installation, yum
3130859	Workaround: Delete the original mlnx-ofed-all-5.* package and recreate the repository with: createrepo RPMS/
	Discovered in Release: 5.8- 1.0.1.1
3149387	 Description: The package neohost-backend (included in MLNX_OFED) has a strict dependency on Python 2.7 and on the existance 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 optionwith-neohost-backend, this sanity check will be overriden 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.
	Keywords: Installation, Python 2
	Workaround: If neohost-backend is needed on a newer system, install Python 2 in advance and create the symbolic link /usr/bin/python -> python2.
	Discovered in Release: 5.8- 1.0.1.1
	Description: 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.
3213777	Keywords: Installation, Oracle Enterprise Linux v9.0
	Workaround: N/A
	Discovered in Release: 5.8- 1.0.1.1
3229904	Description: Restart driver failes to load OFED modules after installing OFED on SLES15sp4 with errata kernel 5.14.21-150400.24.21-default.
	Keywords: Installation

Internal Ref. Number	Issue
	Workaround: Install OFED withadd-kernel-support flag.
	Discovered in Release: 5.8- 1.0.1.1
	Description: VLAN naming is limited to 16 characters (like all other interface names). For names longer than 16 charachters, the kernel generates its own interface name VLAN (VID).
3189424	Keywords: Core, VLAN, Interface Name
	Workaround: Select a name which complies to the 16-characters limitation.
	Discovered in Release: 5.8- 1.0.1.1
	Description: Creating external SFs on BF ARM when the host (x86) operating system does not support SFs may cause the host to crash.
3220855	Keywords: Core, Scalable Functions
	Workaround: N/A
	Discovered in Release: 5.8- 1.0.1.1
3239291	Description: I n some topologies, like logical partitions, mlxfwreset is not supported.
	Keywords: Core, mlxfwreset
	Workaround: N/A
	Discovered in Release: 5.8- 1.0.1.1

Internal Ref. Number	Issue
3114823	Description: The first attempt to create a new iSER connection fails with the following messages in dmesg:
	iSCSI Login timeout on Network Portal <iser_target_ip_addr>:3260 isert: isert_get_login_rx: isert_conn 00000000e9239d52 interrupted before got login req</iser_target_ip_addr>
	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.
	 <u>kernel.org</u> Kernel 5.18 RHEL 9.0 RHEL 8.6

Internal Ref. Number	Issue
	Ubuntu 22.04SLES 15 SP4
	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").
	Workaround: 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.
	Keywords: iSER Initiator
	Discovered in Release: 5.7-1.0.2.0
	Description: 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).
3096911	Workaround: Installing chkconfig before OFED.
	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
3100544	Description: 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 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.
	Workaround: dnf install kernel-modules- # in case it is not the newest.
	Keywords: Installation, RHEL9.x
	Discovered in Release: 5.7-1.0.2.0
3132158	Description: Building rdma-core package on Rocky 8.6 OS caused failure in OFED build.
	Workaround: N/A

Internal Ref. Number	Issue
	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
	Description: Python package is missing, need to install it manually.
2127440	Workaround: Install Python before starting the build.
3137440	Keywords: Installation, Python
	Discovered in Release: 5.7-1.0.2.0
	Description: kernel-macros package does not support building with KMP enabled. KMP needs to be disabled.
3141506	Workaround: Build and install MOFED with KMP disabled (withoutkmp flag).
	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
	Description: kernel-macros package does not support building with KMP enabled. KMP needs to be disabled.
3141506	Workaround: Build and install MOFED with KMP disabled (withoutkmp flag).
	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
	Description: Kernel module packaging is not supported in CtyunOS.
2120027	Workaround: N/A
3129627	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
2971708	 Description: 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. The sysfs interface for roce-enable/disable will be removed entirely for these OSs in a future release. To determine if Devlink can be used to enable or disable RoCE, execute the following console command after starting OFED:

Internal Ref. Number	Issue
	devlink dev param show grep roce
	Devlink supports roce enable/disable if the following line is reflected in the output:
	name enable_roce type generic
	For OSs which do not allow enabling/disabling RoCE via Devlink, the sysfs interface behaves as in the previous 2 releases:
	 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. 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. To determine if Devlink dev reload is supported, execute the following console command (using the bash shell):
	Reload is supported if the output is:
	devlink dev reload DEV [netns { PID NAME ID }]
	Workaround: N/A
	Keywords: Enabling/Disabling RoCE
	Discovered in Release: 5.7-1.0.2.0

Internal Ref. Number	Issue
2998194	 Description: 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).
	Workaround: N/A
	Keywords: NetDev, RHEL, Virtual Functions
	Discovered in Release: 5.6-1.0.3.3
	Description: Rebooting the host while the Arm is down may block the shutdown flow till the Arm is up.
3045436	Workaround: Restart the driver on the host side before reboot.
	Keywords: Reboot, Arm
	Discovered in Release: 5.6-1.0.3.3
	Description:
3040350	 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. When offload is enabled, the physical port must be attached to a bridge.
	Workaround:
	 When removing a physical port from an ovs-dpdk bridge while offload is enabled, need to restart openvswitch after reattaching it. Attach physical port to a bridge according to the desired topology.
	Keywords: OVS-DPDK, Bridge, Offload
	Discovered in Release: 5.6-1.0.3.3
2973726	Description: dec_ttl only work with ConnectX-6. It does not work with ConnectX-5.
	Workaround: N/A

Internal Ref. Number	Issue
	Keywords: OVS-DPDK, dec_ttl
	Discovered in Release: 5.6-1.0.3.3
2946873	Description: Moving to switchdev mode while deleting namespace may cause a deadlock.
	Workaround: Unload mlx5_ib module before moving to Switchdev mode.
	Keywords: ASAP ² , Switchdev, Namespace
	Discovered in Release: 5.6-1.0.3.3
	Description: 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.
2811957	Workaround: N/A
	Keywords: Installation, mlx5_core
	Discovered in Release: 5.6-1.0.3.3
2979243	Description: 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 ./mlnxofedinstalladd-kernelsupport to rebuild the kernel modules for the Centos kernel.
	Workaround: Use add-kernel-support.
	Keywords: Installation,CentOS
	Discovered in Release: 5.6-1.0.3.3
3011440	Description: 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.
	Workaround: Install also mstflint of the exact same version (e.g., apt install mlnx-ofed-all-exact mstflint=4.16.0-1.56xxxx).
	Keywords: Installation,Debian, Ubuntu, MST
	Discovered in Release: 5.6-1.0.3.3

Internal Ref. Number	Issue
3024520	Description: The optioncopy-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 removednot 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.
	Workaround: Remove the files, if needed.
	Keywords: Installation
	Discovered in Release: 5.6-1.0.3.3
	Description: 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.
	Workaround: Either disable building xpmem by addingwithout-xpmem to the command line, or edit the kernel Makefile to make it unset LDFLAGS:
3046601	sed -i -e '/^export ARCH/iLDFLAGS:=' /lib/modules/\$(uname - r)/Makefile
	Note: The Makefile may be located elsewhere, such as the top-level directory of the kernel source directory.
	Keywords: Installation, SLES
	Discovered in Release: 5.6-1.0.3.3
	Description: 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.
3048411	Workaround: Run the command 'dracut -f' to update the initramfs.
	Keywords: Installation
	Discovered in Release: 5.6-1.0.3.3
3048444	Description: OFED installation failed using yum foradd-kernel-support option (building packages without KMP enabled) if libfabric package is installed.

Internal Ref. Number	Issue
	Workaround: Remove libfabric package before OFED installation or use installation script.
	Keywords: Installation, RHEL 8.5
	Discovered in Release: 5.6-1.0.3.3
	Description: OVS topology where the tunnel device is over a VF and the VF representor is connected to a bond is not supported.
3015210	Workaround: N/A
	Keywords: ASAP ² , Tunnel Over VF, LAG, Connection Tracking
	Discovered in Release: 5.6-1.0.3.3
	Description: OVS metering is not support over kernel 5.17.
2020200	Workaround: N/A
3028300	Keywords: ASAP ² ,OVS, Meter, Kernel 5.17
	Discovered in Release: 5.6-1.0.3.3
	Description: Destroying mlxdevm group while SF is attached to it is not supported.
2044255	Workaround: N/A
3044255	Keywords: ASAP ² , mlxdevm, QoS, Group, Scalable Functions, ConnectX-6 Dx
	Discovered in Release: 5.6-1.0.3.3
2900346	Description: 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).
	Workaround: N/A
	Keywords: BlueField-2, Ubuntu, ARP Flux
	Discovered in Release: 5.6-1.0.3.3
3046456	Description: Switching between SwitchDev mode and legacy mode quickly on BlueField-2 can prevent the driver from loading successfully and breaks its health recovery.

Internal Ref. Number	Issue
	Workaround: Pause 60 seconds between state-altering commands to guarantee the driver health recovery is completed successfully.
	Keywords: ASAP ² , Health Recovery
	Discovered in Release: 5.6-1.0.3.3
	Description: Adding vDPA ports over ConnectX-5 devices in ovs-dpdk is not supported and will cause a crash.
2934149	Workaround: N/A
	Keywords: OVS-DPDK, ConnectX-5
	Discovered in Release: 5.6-1.0.3.3
	Description: 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.
2934833	Workaround: N/A
	Keywords: RDMA, Port Toggle
	Discovered in Release: 5.6-1.0.3.3
2901514	Description: Relaxed Ordering is not working properly on Virtual Functions.
	Workaround: N/A
	Keywords: Relaxed Ordering, VF
	Discovered in Release: 5.6-1.0.3.3

Internal Ref. Number	Issue
2870299	Description: Managing SFs is possible using the iproute2 with mlxdevm tool only.
	Workaround: N/A
	Keywords: Scalable Functions
	Discovered in Release: 5.5-1.0.3.2
2869722	Description: OFED packages were built with DKMS disabled since building OFED with DKMS failed due to a problem in the DKMS package

Internal Ref. Number	Issue
	on UOSdkms flag should not be used.
	Workaround: N/A
	Keywords: Installation, DKMS
	Discovered in Release: 5.5-1.0.3.2
	Description: On UOS, IPoIB PKEY may require manual bring up after driver restart.
2870367	Workaround: N/A
	Keywords: Installation, IPoIB, PKEY
	Discovered in Release: 5.5-1.0.3.2
2836032	Description: 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 v5.4x.x.x.
	Workaround: N/A
	Keywords: Software Steering
	Discovered in Release: 5.5-1.0.3.2
	Description: Enabling ARFS in legacy mode and then moving to switchdev mode is not supported and may cause unwanted behavior.
2851639	Workaround: N/A
	Keywords: NetDev, ARFS
	Discovered in Release: 5.5-1.0.3.2
2851639	Description: nvme and iser are not enabled on UOS ARM, because of missing UOS kernel support.
	Workaround: N/A
	Keywords: nvme, iser, UOS ARM
	Discovered in Release: 5.5-1.0.3.2
2860855	Description: Building OFED on RHEL 8.4 with kmp disabled and then installing with yum fails due to some conflicting packages.

Internal Ref. Number	Issue
	Workaround: Remove libfabric and librpmem packages before OFED installation, or addallowerasing option to the installation command.
	Keywords: Installation, RHEL 8.4, kmp, yum
	Discovered in Release: 5.5-1.0.3.2
2865983	Description: OFED packages were built with kmp disabled. Building with kmp enabled fails due to missing packages.
	Workaround: N/A
	Keywords: Installation, kmp
	Discovered in Release: 5.5-1.0.3.2

Internal Ref. Number	Issue
	Description: Only match on lower 32 bit of ct_label is supported.
2050044	Workaround: N/A
2658644	Keywords: ASAP ² , Connection Tracking
	Discovered in Release: 5.4-1.0.3.0
	Description: 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.
2706345	Workaround: To use large number of VFs, set PF_NUM_PF_MSIX to a smaller value of around 32.
	Keywords: ASAP ² ,VF, PF_NUM_PF_MSIX
	Discovered in Release: 5.4-1.0.3.0
2836997	Description: 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).

Internal Ref. Number	Issue
	Workaround: N/A
	Keywords: ASAP ² ,Meter Threshold
	Discovered in Release: 5.4-1.0.3.0
2863456	Description: SA limit by packet count (hard and soft) are supported only on traffic originated from the ECPF. Trying to configure them on VF 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.
	Workaround: N/A
	Keywords: ASAP ² , IPsec Full Offload
	Discovered in Release: 5.4-0.5.1.1
	Description: 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.
2657392	Workaround: N/A
	Keywords: Installation, RHEL, CIFS
	Discovered in Release: 5.4-0.5.1.1
	Description: OpenMPI does not support running across different operating systems and/or CPU architectures.
2800993	Workaround: N/A
	Keywords: OpenMPI
	Description: O pen vSwitch is not supported on the latest operating systems containing only Python3 support.
2399503	Workaround: N/A
	Keywords: Python, O pen vSwitch
2657392	Description: 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.
	Workaround: N/A
	Keywords: Installation, RHEL8.4, CIFS

Internal Ref. Number	Issue
	Discovered in Release: 5.4-0.5.1.1
2782406	Description: Running yum update will upgrade kylin-release to a higher version. The version of this package is used for kylin10sp2 detection so the script will detect kylin 10 instead of kylin10sp2 and use its repository by mistake.
	Workaround: Because there are no special cases for kylin10sp2, the repository that was detected with addingadd-kernel-support to the installation command can be used.
	Keywords: Upgrade, kylin
	Discovered in Release: 5.4-3.0.3.0
	Description: 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.
2755632	Workaround: N/A
	Keywords: SR-IOV, VST, QinQ
	Discovered in Release: 5.4-3.0.3.0
	Description: Non-default MTU (>1500) is not supported with IPsec crypto offload and may cause packet drops.
2780436	Workaround: N/A
	Keywords: IPsec, Crypto Offload, MTU
	Discovered in Release: 5.4-3.0.3.0
2726021	 Description: 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. On kylin OS, the version of kernel-rpm-macros package does not support -p option needed to support kmp, so it will stay disabled.
	Workaround: N/A
	Keywords: Installation, openEuler
	Discovered in Release: 5.4-3.0.3.0

Internal Ref. Number	Issue
2750653	Description: Running fragmented traffic in RHEL 8.3 (4.18.0-240.el8.x86_64) may cause call trace in build_skb.
	Workaround: Update to RHEL 8.3 z-stream 4.18.0- 240.22.1.el8_3.x86_64.
	Keywords: RHEL 8.3, Kernel Panic, Call Trace, fr
	Discovered in Release: 5.4-1.0.3.0
	Description: 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.
2629375	Workaround: N/A
	Keywords: ASAP ² , Connection Tracking, Label
	Discovered in Release: 5.4-1.0.3.0
	Description: Installation in the package manager mode under SLES 15.x may require user-intervention if the original libibverbs is installed.
2707997	Workaround: zypper installforce-resolution mlnx-ofed-all
	Keywords: Installation, libibverbs
	Discovered in Release: 5.4-1.0.3.0
	Description: Installation in the package manager mode under SLES 15.x may require user-intervention if the original libopenvswitch is installed.
2708531	Workaround: zypper installforce-resolution mlnx-ofed-all
	Keywords: Installation
	Discovered in Release: 5.4-1.0.3.0
	Description: Congested TCP lock for kTLS TX device offload traffic compromises the performance.
2703043	Workaround: Disable TCP selective acknowledgement: echo 0 > /proc/sys/net/ipv4/tcp_sack
	Keywords: kTLS TX
	Discovered in Release: 5.4-1.0.3.0
2676405	Description: If the package interface-rename is active (on XenServer, for example), the interface renaming by the OFED will not be done to eliminate conflicts.

Internal Ref. Number	Issue
	Workaround: N/A
	Keywords: Interface Renaming
	Discovered in Release: 5.4-1.0.3.0
	Description: Offload of rules which redirect from VF on one PF to VF on second PF is not supported on socket-direct devices.
2687943	Workaround: N/A
	Keywords: ASAP ² , Socket-Direct
	Discovered in Release: 5.4-1.0.3.0
	Description: 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.
2678672	Workaround: N/A
	Keywords: SwitchDev, mlx5, Tunnel Traffic
	Discovered in Release: 5.4-1.0.3.0
	Description: On PPC systems when EEH is enabled, running fw sync reset (either by mlxfwreset with flagsync 1 or by devlink dev reload action fw_activate), the EEHmay 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.
2566548	Workaround: Administrator may disable EEH before running firmware sync reset on the device.
	Keywords: PPC, EEH
	Discovered in Release: 5.4-1.0.3.0
2617950	Description: 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.
	Workaround: N/A
	Keywords: Ethtool
	Discovered in Release: 5.4-1.0.3.0
2390731	Description: Ethtool does not display Port Speed advertised/capability above 100Gb/s over and below kernels 5.0, even when supported.

Internal Ref. Number	Issue
	Workaround: N/A
	Keywords: Ethtool, Port Speed
	Discovered in Release: 5.4-1.0.3.0

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

Internal Ref. Number	Issue
	Description: 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.
2585575	Workaround: N/A
	Keywords: Firmware Upgrade
	Discovered in Release: 5.3-1.0.0.1
	Description: Conducting a firmware reset or unbinding the PF while in switchdev mode may cause a kernel crash.
2582565	Workaround: N/A
	Keywords: SwitchDev, ASAP ² , Unbind, Firmware Reset
	Discovered in Release: 5.3-1.0.0.1
2587802	Description: PTP synchronization may be lost while using tx_port_ts private flag.
	Workaround: Toggle private flag: ethtoolset-priv-flags <ifs> tx_port_ts off ethtoolset-priv-flags <ifs> tx_port_ts on restart ptp4l application</ifs></ifs>

Internal Ref. Number	Issue
	Keywords: PTP Synchronization
	Discovered in Release: 5.3-1.0.0.1
	Description: When running kernel 5.8 and bellow or RHEL 8.2 and below, sampled packets do not support tunnel information.
2574943	Workaround: N/A
	Keywords: ASAP ² , sFLOW
	Discovered in Release: 5.3-1.0.0.1
2568417	Description: 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.
	Workaround: N/A
	Keywords: Upgrade; mlx5_fpga_tools
	Discovered in Release: 5.3-1.0.0.1
2506425	Description: 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.
2000-20	Workaround: N/A
	Keywords: Installation; modules; kmod
	Discovered in Release: 5.3-1.0.0.1
2492509	Description: When installing the driver on OpenEuler or on EulerOS 2.0SP9, rebuilding the drivers (add-kernel-support) with thekmp 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.
	Workaround: N/A
	Keywords: add-kernel-support
	Discovered in Release: 5.3-1.0.0.1

Internal Ref. Number	Issue
	Description: 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.
2479327	Workaround: N/A
	Keywords: Upgrade; SLES 12; add-kernel-support
	Discovered in Release: 5.3-1.0.0.1
	Description: 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.
2584441	Workaround: N/A
	Keywords: Upgrade; SLES 12; add-kernel-support
	Discovered in Release: 5.3-1.0.0.1
	Description: When setting MTU to low values, such as 68 bytes, packets may fail on oversize.
2460865	Workaround: N/A
	Keywords: MTU
	Discovered in Release: 5.3-1.0.0.1
	Description: On kernels based on RedHat 7.2, the "tx_port_ts" feature, as set by ethtool —set-priv-flags, is disabled.
2383318	Workaround: N/A
	Keywords: RedHat; tx_port_ts
	Discovered in Release: 5.3-1.0.0.1
	Description: An OvS-DPDK crash might occur while doing live-migration for VMs that use virtio-interfaces that are accelerated using OvS-DPDK vDPA ports.
2575647	Workaround: N/A
	Keywords: OvS-DPDK vDPA, Live-migration
	Discovered in Release: 5.3-1.0.0.1

Internal Ref. Number	Issue
	Description: 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.
2430071	Workaround: N/A
	Keywords: IPoIB devlink reload
	Discovered in Release: 5.2-2.2.0.0
2202700	Description: 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.
2302786	Workaround: Install MLNX_OFED withadd-kernel-support.
	Keywords: EulerOS; kABI; installation;add-kernel-support
	Discovered in Release: 5.2-1.0.4.0
	Description: A crash in the TLS Rx socket cleanup flow may occur due to a kernel issue where a wrong extra call to tls_dev_del is made.
2398281	Workaround: N/A
	Keywords: TLS RX device offload
	Discovered in Release: 5.2-1.0.4.0
	Description: 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.
2407415	Workaround: Install MLNX_OFED withadd-kernel-support.
	Keywords: OpenEuler; Aarch64; installation;add-kernel-support
	Discovered in Release: 5.2-1.0.4.0
2348077	Description: RDMA device name for VFs may change after resetting all VFs at once.
	Workaround: 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.

Internal Ref. Number	Issue
	Keywords: RDMA; VF
	Discovered in Release: 5.2-1.0.4.0
	Description: 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.
2381713	Workaround: Either install MLNX_OFED usingadd-kernel-support, or rebuild the non-default custom kernel to include these modules.
	Keywords: esp4_offload; esp6_offload; kernel, Debian
	Discovered in Release: 5.2-1.0.4.0
	Description : 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.
2382898	Workaround: N/A
	Keywords: GENEVE; stag; VLAN; UDP
	Discovered in Release: 5.2-1.0.4.0
	Description: 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 <gid> error=-28</gid>
2326155	Workaround: N/A
	Keywords: RoCE;ib_cache_gid_add
	Discovered in Release: 5.2-1.0.4.0
	Description: Running XDP over an IP tunnel may fail when working with kernels as old as version 4.14.
2329654	Workaround: N/A
	Keywords: XDP, Kernel
	Discovered in Release: 5.2-1.0.4.0
2249156	Description: MLNX_OFED installation will remove qperf package in case it was done after qperf installation.
	Workaround : Make sure to install qperf package after installing MLNX_OFED, or re-install qperf after installing MLNX_OFED.
	Keywords: Installation; qperf

Internal Ref. Number	Issue
	Discovered in Release: 5.2-1.0.4.0
	Description: OFED installation requires kernel config CONFIG_DEBUG_INFO to be set.
2355956	Workaround: N/A
	Keywords: Installation; CONFIG_DEBUG_INFO
	Discovered in Release: 5.2-1.0.4.0
	Description: 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.
2362781	Workaround: N/A
	Keywords: Openibd; Inbox; Ubuntu; mlx5_ib
	Discovered in Release: 5.2-1.0.4.0
2367659	Description: 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.×8664/extra/iser/ib_iser.ko needs unknown symbol ib_fmr_pool_unmap depmod: WARNING: /lib/modules/4.18.0-240.el8.×8664/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.
	Workaround: N/A
	Keywords: YUM; RPM; symbol; depmod; ISER; SRP
	Discovered in Release: 5.2-1.0.4.0
2385269	Description: The number of connections offloaded is limited to 100K when working with Kernel v5.9.
	Workaround: N/A
	Keywords: ASAP2; Connection Tracking; Kernel
	Discovered in Release: 5.2-1.0.4.0
2393169	Description: Mirroring is not supported with Connection Tracking when the source port is a VxLAN device.

Internal Ref. Number	Issue
	Workaround: N/A
	Keywords: ASAP ² ; Connection Tracking; Mirroring
	Discovered in Release: 5.2-1.0.4.0
	Description: 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.
2395082	Workaround: N/A
	Keywords: ASAP ² ;SwitchDev; call trace; kernel; tcf_block_unbind
	Discovered in Release: 5.2-1.0.4.0

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

Internal Ref. Number	Issue
2295735	Description: Upgrading from legacy (mlnx-libs) to the current rdma-core based build using the apt-get (package manager) fails.
	Workaround: To perform this upgrade, either use the installer script or uninstall the old packages and install the new packages.
	Keywords: Legacy, mlnx-libs, rdma-core, apt, apt-get, installation
	Discovered in Release: 5.1-2.5.8.0
	Description: Downgrading the firmware version for ConnectX-6 cards using "mlnx_ofed_installfw-update-onlyforce-fw-update" fails.
2248996	Workaround: Manually downgrade the firmware version - please see <u>Firmware Update Instructions</u> .
	Keywords: Firmware, ConnectX-6
	Discovered in Release: 5.1-0.6.6.0
	Description: 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.
2175930	Workaround: Update the ethtool utility to version 5.6 on such systems in order to avoid the call trace.
	Keywords: PPC, ethtool_notify, kernel
	Discovered in Release: 5.1-0.6.6.0
2198764	Description: 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.
	Workaround: Currently, openibd is a sysv-init script that you can enable manually by running: update-rc.d openibd defaults
	Keywords: chroot, Debian , Ubuntu, openibd
	Discovered in Release: 5.1-0.6.6.0
2237134	Description: Running connection tracking (CT) with FW steering may cause CREATE_FLOW_TABLE command to fail with syndrome.
	Workaround: Configure OVS to use a single handler-thread: #ovs-vsctl set Open_vSwitch . other_config:n-handler-threads=1

Internal Ref. Number	Issue
	Keywords: Connection tracking, ASAP, OVS, FW steering
	Discovered in Release: 5.1-0.6.6.0
2239894	Description: Running OpenVSwitch offload with high traffic throughput can cause low insertion rate due to high CPU usage.
	Workaround: Reduce the number of combined channels of the uplink using "ethtool -L".
	Keywords: Insertion rate, ASAP2
	Discovered in Release: 5.1-0.6.6.0
	Description: Header rewrite action is not supported over RHEL/CentOS 7.4.
2240671	Workaround: N/A
	Keywords: ASAP, header rewrite, RHEL, RedHat, CentOS, OS
	Discovered in Release: 5.1-0.6.6.0
	Description: Tunnel offload (encap/decap) may cause kernel panic if nf_tables module is not probed.
2242546	Workaround: Make sure to probe the nf_tables module before inserting any rule.
	Keywords: Kernel v5.7, ASAP, kernel panic
	Discovered in Release: 5.1-0.6.6.0
	Description: IPsec packets are dropped during heavy traffic due to a bug in net/xfrm Linux Kernel.
2143007	Workaround: Make sure the Kernel is modified to apply the following patch: "xfrm: Fix double ESP trailer insertion in IPsec crypto offload".
	Keywords: IPsec, xfrm
	Discovered in Release: 5.1-0.6.6.0
2225952	Description: VF mirroring with TC policy skip_sw is not supported on RHEL/CentOS 7.4, 7.5 and 7.6 OSs.
	Workaround: N/A
	Keywords: ASAP ² , Mirroring, RHEL, RedHat, OS
	Discovered in Release: 5.1-0.6.6.0

Internal Ref. Number	Issue
2216521	Description: After upgrading MLNX_OFED from v5.0 or earlier, ibdev2netdev utility changes the installation prefix to /usr/sbin. Therefore, it cannot be found while found in the same SHELL environment.
	Workaround: After installing MLNX_OFED, log out and log in again to refresh the SHELL environment.
	Keywords: ibdev2netdev
	Discovered in Release: 5.1-0.6.6.0
	Description: Rules with VLAN push/pop, encap/decap and header rewrite actions together are not supported.
0000500	Workaround: N/A
2202520	Keywords: ASAP ² , SwitchDev, VLAN push/pop, encap/decap, header rewrite
	Discovered in Release: 5.1-0.6.6.0
2210752	Description: Switching from Legacy mode to SwitchDev mode and vice-versa while TC rules exist on the NIC will result in failure.
	Workaround: Before attempting to switch mode, make sure to delete all TC rules on the NIC or stop OpenvSwitch.
	Keywords: ASAP ² , Devlink, Legacy SR-IOV
	Discovered in Release: 5.1-0.6.6.0

Internal Ref. Number	Issue
2125036/21 25031	Description: 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.
	Workaround: Make sure to uninstall and re-install MLNX_OFED to complete the upgrade.
	Keywords: Installation, UPSTREAM_LIBS, MLNX_LIBS
	Discovered in Release: 5.0-2.1.8.0
2105447	Description: hns_roce warning messages will appear in the dmesg after reboot on Euler2 SP3 OSs.

Internal Ref. Number	Issue
	Workaround: N/A
	Keywords: hns_roce, dmesg, Euler
	Discovered in Release: 5.0-2.1.8.0
	Description: Multiple driver restarts may cause IPoIB soft lockup.
2110221	Workaround: N/A
2110321	Keywords: Driver restart, IPoIB
	Discovered in Release: 5.0-2.1.8.0
	Description: 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.
2112251	Workaround: Define geneve tunnel's remote endpoint using IPv4.
	Keywords: Kernel, Geneve, IPv4, IPv6, MTU, fragmentation
	Discovered in Release: 5.0-2.1.8.0
	Description: 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.
2102902	Workaround: N/A
	Keywords: TLS offload, mlx5e
	Discovered in Release: 5.0-2.1.8.0
2111534	Description: 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.
	Workaround: N/A
	Keywords: TLS offload, mlx5e
	Discovered in Release: 5.0-2.1.8.0

Internal Ref. Number	Issue
2035950	Description: 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.
	 Removing PF from the bond (using ifdown, ip link or any other function) Attempting to disable SR-IOV
	Workaround: N/A
	Keywords: VF LAG, binding, firmware, FW, PF, SR-IOV
	Discovered in Release: 5.0-1.0.0.0
	Description: When working with OSs with Kernel v4.10, bonding module does not allow setting MTUs larger than 1500 on a bonding interface.
2044544	Workaround: Upgrade your Kernel version to v4.11 or above.
	Keywords: Bonding, MTU, Kernel
	Discovered in Release: 5.0-1.0.0.0
	Description: Libibverbs dependencies are removed during OFED installation, requiring manual installation of libraries that OFED does not reinstall.
1882932	Workaround: Manually install missing packages.
	Keywords: libibverbs, installation
	Discovered in Release: 5.0-1.0.0.0
2058535	Description: ibdev2netdev command returns duplicate devices with different ports in SwitchDev mode.
	Workaround: Use /opt/mellanox/iproute2/sbin/rdma link show command instead.
	Keywords: ibdev2netdev
	Discovered in Release: 5.0-1.0.0.0
2072568	Description: In RHEL/CentOS 7.2 OSs, adding drop rules when act_gact is not loaded may cause a kernel crash.
	Workaround: Preload all needed modules to avoid such a scenario (cls_flower, act_mirred, act_gact, act_tunnel_key and act_vlan).

Internal Ref. Number	Issue
	Keywords: RHEL/CentOS 7.2, Kernel 4.9, call trace, ASAP
	Discovered in Release: 5.0-1.0.0.0
	Description: VF LAG configuration is not supported when the NUM_OF_VFS configured in mlxconfig is higher than 64.
2093698	Workaround: N/A
	Keywords: VF LAG, SwitchDev mode, ASAP
	Discovered in Release: 5.0-1.0.0.0
	Description: Devlink health dumps are not supported on kernels lower than v5.3.
2093746	Workaround: N/A
	Keywords: Devlink, health report, dump
	Discovered in Release: 5.0-1.0.0.0
	Description: Sending packets larger than MTU is not supported when working with OVS-DPDK.
2000590	Workaround: N/A
	Keywords: MTU, OVS-DPDK
	Discovered in Release: 5.0-1.0.0.0
	Description: Moving VF from SwitchDev mode to Legacy mode while the representor is being used by OVS-DPDK results in a segmentation fault.
2062900	Workaround: To move VF to Legacy mode with no error, make sure to delete the ports from the OVS.
	Keywords: SwitchDev, Legacy, representor, OVS-DPDK
	Discovered in Release: 5.0-1.0.0.0
2075942	Description: Huge pages configuration is lost each time the server is configured.
	Workaround: Re-configure the huge pages after each reboot, or configure them as a kernel parameter.
	Keywords: Huge pages, reboot, OVS-DPDK
	Discovered in Release: 5.0-1.0.0.0

Internal Ref. Number	Issue
2067012	Description: MLNX_OFED cannot be installed on Debian 9.11 OS in SwitchDev mode.
	Workaround: Install OFED with the flagadd-kernel-support.
	Keywords: ASAP, SwitchDev, Debian, Kernel
	Discovered in Release: 5.0-1.0.0.0
2036572	Description: 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.
	Workaround: N/A
	Keywords: rdma-core, write-combining, MMIO buffer
	Discovered in Release: 5.0-1.0.0.0

Internal Ref. Number	Issue
	Description: 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.
-	Workaround: Install the argparse module manually.
	Keywords: Python, MFT, argparse, installation
	Discovered in Release: 4.7-3.2.9.0
1997230	Description: Running mlxfwreset or unloading mlx5_core module while contrak flows are offloaded may cause a call trace in the kernel.
	Workaround: Stop OVS service before calling mlxfwreset or unloading mlx5_core module.
	Keywords: Contrak, ASAP, OVS, mlxfwrest, unload
	Discovered in Release: 4.7-3.2.9.0
1955352	Description : Moving 2 ports to SwitchDev mode in parallel is not supported.
	Workaround: N/A
	Keywords: ASAP, SwitchDev

Internal Ref. Number	Issue
	Discovered in Release: 4.7-3.2.9.0
	Description: VxLAN IPv6 offload is not supported over CentOS/RHEL v7.2 OSs.
1979958	Workaround: N/A
	Keywords: Tunnel, VXLAN, ASAP, IPv6
	Discovered in Release: 4.7-3.2.9.0
	Description: PRIO_TAG_REQUIRED_EN configuration is not supported and may cause call trace.
1991710	Workaround: N/A
	Keywords: ASAP, PRIO_TAG, mstconfig
	Discovered in Release: 4.7-3.2.9.0
	Description: Enabling ECMP offload requires the VFs to be unbound and VMs to be shut down.
1967866	Workaround: N/A
	Keywords: ECMP, Multipath, ASAP ²
	Discovered in Release: 4.7-3.2.9.0
	Description: On Ubuntu, Debian and RedHat 8 and above OSS, parsing the mfa2 file using the mstarchive might result in a segmentation fault.
1021001	Workaround: Use mlxarchive to parse the mfa2 file instead.
1921981	Keywords: MFT, mfa2, mstarchive, mlxarchive, Ubuntu, Debian, RedHat, operating system
	Discovered in Release: 4.7-1.0.0.1
1840288	Description: MLNX_OFED does not support XDP features on RedHat 7 OS, despite the declared support by RedHat.
	Workaround: N/A
	Keywords: XDP, RedHat
	Discovered in Release: 4.7-1.0.0.1
1821235	Description: 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.

Internal Ref. Number	Issue
	Workaround: Use the both actions at the end of the flow process.
	Keywords: Flow steering
	Discovered in Release: 4.7-1.0.0.1

Internal Ref. Number	Issue	
1504785	Description: A lost interrupt issue in pass-through virtual machines may prevent the driver from loading, followed by printing managed pages errors to the dmesg.	
	Workaround: Restart the driver.	
	Keywords: VM, virtual machine	
	Discovered in Release: 4.6-1.0.1.1	
	Description: Unbinding PFs on LAG devices results in a "Failed to modify QP to RESET" error message.	
1764415	Workaround: N/A	
	Keywords: RoCE LAG, unbind, PF, RDMA	
	Discovered in Release: 4.6-1.0.1.1	
1806565	Description: 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.	
	Workaround: Use RoCEv2 default GID.	
	Keywords: RoCE	
	Discovered in Release: 4.6-1.0.1.1	
	Description: Aging is not functional on bond device in RHEL 7.6.	
-	Workaround: N/A	
	Keywords: VF LAG, ASAP ²	
	Discovered in Release: 4.6-1.0.1.1	
1747774	Description: 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	

Internal Ref. Number	Issue
	port 1 and half with port 2. All the traffic on the same ring will be sent from the same port.
	Workaround: N/A
	Keywords: VF LAG, ASAP ²
	Discovered in Release: 4.6-1.0.1.1
	Description: 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
1753629	Workaround: N/A
	Keywords: Bonding, slave
	Discovered in Release: 4.6-1.0.1.1
	Description: Uninstalling MLNX_OFED automatically results in the uninstallation of several libraries that are included in the MLNX_OFED package, such as InfiniBand-related libraries.
1712068	Workaround: If these libraries are required, reinstall them using the local package manager (yum/dnf).
	Keywords: MLNX_OFED libraries
	Discovered in Release: 4.6-1.0.1.1
	Description: 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.
-	Workaround: N/A
	Keywords: MFT compatible
	Discovered in Release: 4.6-1.0.1.1
	Description: On ConnectX-4 HCAs, GID index for RoCE v2 is inconsistent when toggling between enabled and disabled interface modes.
1730840	Workaround: N/A
	Keywords: RoCE v2, GID
	Discovered in Release: 4.6-1.0.1.1

Internal Ref. Number	Issue
1717428	Description : On kernels 4.10-4.14, MTUs larger than 1500 cannot be set for a GRE interface with any driver (IPv4 or IPv6).
	Workaround: Upgrade your kernel to any version higher than v4.14.
	Keywords: Fedora 27, gretap, ip_gre, ip_tunnel, ip6_gre, ip6_tunnel
	Discovered in Release: 4.6-1.0.1.1
	Description : Driver reload takes several minutes when a large number of VFs exists.
1748343	Workaround: N/A
	Keywords: VF, SR-IOV
	Discovered in Release: 4.6-1.0.1.1
	Description : 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.
1733974	Workaround: N/A
	Keywords: mlx5
	Discovered in Release: 4.6-1.0.1.1
	Description : On ConnectX-6 HCAs and above, an attempt to configure advertisement (any bitmap) will result in advertising the whole capabilities.
-	Workaround: N/A
	Keywords: 200Gb/s, advertisement, Ethtool
	Discovered in Release: 4.6-1.0.1.1

Internal Ref. Number	Issue
1699289	Description: 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.
	Workaround: Make sure to enable HW LRO in the driver: ethtool -k <intf> lro ethtoolset-priv-flag <intf> hw_lro on</intf></intf>

Internal Ref. Number	Issue
	Keywords: HW LRO, ConnectX-5 and above
	Discovered in Release: 4.5-1.0.1.0
	Description : 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.
1403313	Workaround : Make sure to update the Kernel version to v4.15 or above.
	Keywords: VF, PF, IOMMU, Kernel, OS
	Discovered in Release: 4.5-1.0.1.0
_	 SLES12 SP3 SLES12 SP4 SLES15 Fedora 28 RHEL7.1 RHEL7.4 ALT (Pegas 1.0) REL 7.5 RHEL7.6 XenServer 4.9
	Workaround: N/A
	Keywords: NEO-Host, operating systems
	Discovered in Release: 4.5-1.0.1.0
1521877	Description : On SLES 12 SP1 OSs, a kernel tracepoint issue may cause undefined behavior when inserting a kernel module with a wrong parameter.
	Workaround: N/A
	Keywords : mlx5 driver, SLES 12 SP1
	Discovered in Release: 4.5-1.0.1.0

User Manual

- Introduction
- Installation
- Features Overview and Configuration
- <u>Programming</u>
- InfiniBand Fabric Utilities
- <u>Troubleshooting</u>
- Common Abbreviations and Related Documents

Introduction

i) Note

MLNX_OFED has transitioned into DOCA-Host, and now available as DOCA-OFED (learn about DOCA-Host profiles <u>here</u>).

MLNX_OFED last standalone release is October 2024 Long Term Support (3 years). Starting January 2025 all new features will be included in DOCA-OFED only.

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	 InfiniBand: SDR, FDR, EDR, HDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE², 100GbE²
BlueField		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-7		 InfiniBand: EDR, HDR100, HDR, NDR200, NDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE², 100GbE², 200GbE³
ConnectX-6 Lx		• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE ²
ConnectX-6 Dx		 Ethernet: 10GbE, 25GbE, 40GbE, 50GbE², 100GbE², 200GbE²
ConnectX-6		 InfiniBand: SDR, FDR, EDR, HDR Ethernet: 10GbE, 25GbE, 40GbE, 50GbE², 100GbE², 200GbE²
ConnectX- 5/ConnectX-5 Ex		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 100GbE
ConnectX-4 Lx		• Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE
ConnectX-4		 InfiniBand: SDR, QDR, FDR, FDR10, EDR Ethernet: 1GbE, 10GbE, 25GbE, 40GbE, 50GbE, 56GbE¹, 100GbE

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:

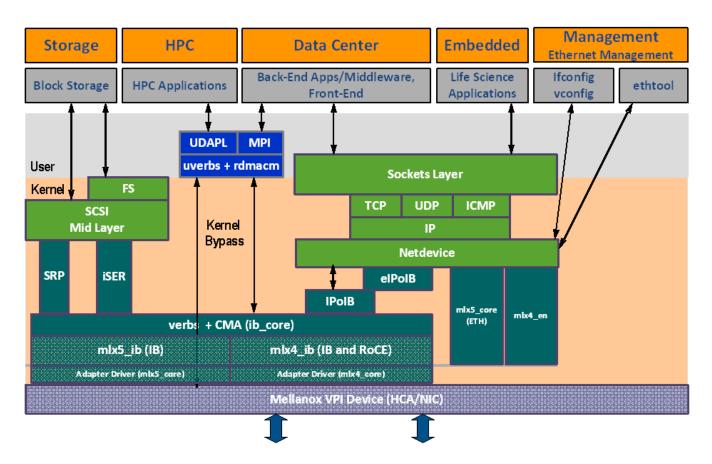
• NVIDIA Messaging Accelerator (VMA[™]) software: Socket acceleration library that performs OS bypass for standard socket-based applications.

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).

- 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



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:

(i) 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 usermode 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 "<u>iSCSI Extensions for RDMA (iSER)</u>" 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 "<u>SRP—SCSI RDMA</u> <u>Protocol</u>" 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) 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 BeBenchmarkand 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 OFED¹.

1. OpenSM is disabled by default. See "<u>NVIDIA SM</u>" 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).

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:

- mlnxofedinstall—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.



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.

	The parameter name for selecting the profile. The supported values for profiles are:	
prof_sel	 0 - for medium resources, medium performance 1 - for low resources 2 - for high performance (int) (default) 3 - fast unload 4 - without DCT enable 	
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_grou ps	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)

ib_ipoib Parameters

max_nonsrq_c onn_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_si ze	Number of descriptors in send queue (int)
recv_queue_siz e	Number of descriptors in receive queue (int)
debug_level	Enable debug tracing if > 0 (int)
ipoib_enhance d	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 the Driver
- Uninstalling the Driver
- Updating Firmware After Installation
- UEFI Secure Boot
- Performance Tuning

Hardware and Software Requirements

Requirement s	Description
Platforms	A server platform with an adapter card based on one of the following NVIDIA' VPI HCA devices listed in <u>Supported NICs Speeds</u> table. For the list of supported architecture platforms, please refer to the <u>NVIDIA OFED Release Notes</u> .
Required Disk Space for Installation	1GB
Device ID	For the latest list of device IDs, please visit NVIDIA website.
Operating System	Linux operating system. For the list of supported operating system distributions and kernels, please refer to the <u>NVIDIA OFED Release Notes</u> .
Installer Privileges	The installation requires administrator (root) privileges on the target machine.

Downloading the Drivers

1. Verify that the system has a NVIDIA network adapter (HCA/NIC) installed.

The following example shows a system with an installed NVIDIA HCA:

```
# lspci -v | grep Mellanox
```

```
86:00.0 Network controller [0207]: Mellanox Technologies
MT27620 Family
    Subsystem: Mellanox Technologies Device 0014
86:00.1 Network controller [0207]: Mellanox Technologies
MT27620 Family
    Subsystem: Mellanox Technologies Device 0014
```

Note: For ConnectX-5 Socket Direct adapters, use ibdev2netdev to display the installed card and the mapping of logical ports to physical ports. Example:

```
[root@gen-l-vrt-203 ~]# ibdev2netdev -v | grep -i MCX556M-
ECAT-S25
0000:84:00.0 mlx5_10 (MT4119 - MCX556M-ECAT-S25SN) CX556M -
ConnectX-5 QSFP28 fw 16.22.0228 port 1 (DOWN ) ==> p2p1 (Down)
0000:84:00.1 mlx5_11 (MT4119 - MCX556M-ECAT-S25SN) CX556M -
ConnectX-5 QSFP28 fw 16.22.0228 port 1 (DOWN ) ==> p2p2 (Down)
0000:05:00.0 mlx5_2 (MT4119 - MCX556M-ECAT-S25SN) CX556M -
ConnectX-5 QSFP28 fw 16.22.0228 port 1 (DOWN ) ==> p5p1 (Down)
0000:05:00.1 mlx5_3 (MT4119 - MCX556M-ECAT-S25SN) CX556M -
ConnectX-5 QSFP28 fw 16.22.0228 port 1 (DOWN ) ==> p5p1 (Down)
```

(i) Note

- Each PCI card of ConnectX-5 Socket Direct has a different PCI address. In the output example above, the first two rows indicate that one card is installed in a PCI slot with PCI Bus address 84 (hexadecimal), and PCI Device Number 00, and PCI Function Number 0 and 1. RoCE assigned mlx5_10 as the logical port, which is the same as netdevice p2p1, and both are mapped to physical port of PCI function 0000:84:00.0.
- RoCE logical port mlx5_2 of the second PCI card (PCI Bus address 05) and netdevice p5p1 are mapped to physical

port of PCI function 0000:05:00.0, which is the same physical port of PCI function 0000:84:00.0.

MT4119 is the PCI Device ID of the ConnectX-5 adapters family.

For more details, please refer to ConnectX-5 Socket Direct Hardware User Manual, available at <u>nvidia.com/en-</u><u>us/networking/</u>.

2. Download the ISO image to your host.

The image's name has the format MLNX_OFED_LINUX-<ver>-<OS label><CPU arch>.iso. It can be download from <u>nvidia.com/en-us/networking/</u> \rightarrow Products \rightarrow Software \rightarrow InfiniBand Drivers \rightarrow MLNX_OFED.

- 1. Scroll down to the Download wizard, and click the Download tab.
- 2. Choose your relevant package depending on your host operating system.
- 3. Click the desired ISO/tgz package.
- 4. To obtain the download link, accept the End User License Agreement (EULA).
- 3. Use the sa256sum utility to confirm the file integrity of your ISO image. Run the following command and compare the result to the value provided on the download page.

host1\$ sha256sum MLNX_OFED_LINUX-<ver>-<OS label>.iso

Installing the Driver

i Note

For installation of MLNX_OFED on Community OSs, please see section "Installation on Community Operating Systems" below.

Installation Script

The installation script, mlnxofedinstall, performs the following:

- Discovers the currently installed kernel
- Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Installs the MLNX_OFED_LINUX binary RPMs (if they are available for the current kernel)
- Identifies the currently installed InfiniBand and Ethernet network adapters and automatically upgrades the firmware

Note: To perform a firmware upgrade using customized firmware binaries, a path can be provided to the folder that contains the firmware binary files, by running --fw-image-dir. Using this option, the firmware version embedded in the MLNX_OFED package will be ignored.

Example:

./mlnxofedinstall --fw-image-dir /tmp/my_fw_bin_files

(i) Note

If the driver detects unsupported cards on the system, it will abort the installation procedure. To avoid this, make sure to add --skipunsupported-devices-check flag during installation.

Usage

./mnt/mlnxofedinstall [OPTIONS]

The installation script removes all previously installed OFED packages and re-installs from scratch. You will be prompted to acknowledge the deletion of the old packages.

(i) Note

Pre-existing configuration files will be saved with the extension ".conf.rpmsave".

- If you need to install OFED on an entire (homogeneous) cluster, a common strategy is to mount the ISO image on one of the cluster nodes and then copy it to a shared file system such as NFS. To install on all the cluster nodes, use cluster-aware tools (suchaspdsh).
- If your kernel version does not match with any of the offered pre-built RPMs, you can add your kernel version by using the "mlnx_add_kernel_support.sh" script located inside the MLNX_OFED package.

(i) Note

On Redhat and SLES distributions with errata kernel installed there is no need to use the mlnx_add_kernel_support.sh script. The regular installation can be performed and weak-updates mechanism will create symbolic links to the MLNX_OFED kernel modules.

i) Note

If you regenerate kernel modules for a custom kernel (using -add-kernel-support), the packages installation will not involve automatic regeneration of the initramfs. In some cases, such as a system with a root filesystem mounted over a ConnectX card, not regenerating the initramfs may even cause the system to fail to reboot.

In such cases, the installer will recommend running the following command to update the initramfs:

dracut -f

On some OSs, dracut -f might result in the following error message which can be safely ignore.

libkmod: kmod_module_new_from_path: kmod_module 'mdev' already exists with different path

The "mlnx_add_kernel_support.sh" script can be executed directly from the mlnxofedinstall script. For further information, please see '--add-kernel-support' option below.



On Ubuntu and Debian distributions drivers installation use Dynamic Kernel Module Support (DKMS) framework. Thus, the drivers' compilation will take place on the host during MLNX_OFED installation.

Therefore, using "mlnx_add_kernel_support.sh" is irrelevant on Ubuntu and Debian distributions.

Example: The following command will create a MLNX_OFED_LINUX ISO image for RedHat 7.3 under the /tmp directory.

./MLNX_OFED_LINUX-x.x-rhel7.3x86_64/mlnx_add_kernel_support.sh -m /tmp/MLNX_OFED_LINUX-

```
x.x-x-rhel7.3-x86_64/ --make-tgz
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.3
under /tmp directory.
All Mellanox, OEM, OFED, or Distribution IB packages will be
removed.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_ofed_iso.21642.log
Building OFED RPMs. Please wait...
Removing OFED RPMs...
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel7.3-x86_64-ext.tgz
```

- The script adds the following lines to /etc/security/limits.conf for the userspace components such as MPI:
 - * soft memlock unlimited
 - * hard memlock unlimited
 - These settings set the amount of memory that can be pinned by a userspace application to unlimited. If desired, tune the value unlimited to a specific amount of RAM.

For your machine to be part of the InfiniBand/VPI fabric, a Subnet Manager must be running on one of the fabric nodes. At this point, OFED for Linux has already installed the OpenSM Subnet Manager on your machine.

For the list of installation options, run: ./mlnxofedinstall --h

Installation Procedure

This section describes the installation procedure of MLNX_OFED on NVIDIA adapter cards.

- 1. Log in to the installation machine as root.
- 2. Mount the ISO image on your machine.

```
host1# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU
arch>.iso /mnt
```

3. Run the installation script.

```
/mnt/mlnxofedinstall
Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs
This program will install the MLNX_OFED_LINUX package on your
machine.
Note that all other Mellanox, OEM, OFED, RDMA or Distribution
IB packages will be removed.
Those packages are removed due to conflicts with
MLNX_OFED_LINUX, do not reinstall them.
Starting MLNX_OFED_LINUX-x.x.x installation ...
.....
Installation finished successfully.
Attempting to perform Firmware update...
Querying Mellanox devices firmware ...
```

(i) Note

For unattended installation, use the --force installation option while running the MLNX_OFED installation script:

/mnt/mlnxofedinstall --force

(i) Note

MLNX_OFED for Ubuntu should be installed with the following flags in chroot environment:

./mlnxofedinstall --without-dkms --add-kernel-support --kernel -without-fw-update --force

For example:

./mlnxofedinstall --without-dkms --add-kernel-support --kernel 3.13.0-85-generic --without-fw-update --force

Note that the path to kernel sources (--kernel-sources) should be added if the sources are not in their default location.

i) Note

In case your machine has the latest firmware, no firmware update will occur and the installation script will print at the end of installation a message similar to the following:

Device #1:

Device Type: ConnectX4

Part Number: MCX456A-ECA

Description: ConnectX-4 VPI adapter card; EDR IB (100Gb/s) and 100GbE; dual-port QSFP28; PCIe3.0 x16; ROHS R6

PSID: MT_2190110032

PCI Device Name: 0b:00.0

Base MAC: 0000e41d2d5cf810

Versions: Current Available

FW 12.14.0114 12.14.0114

(i) Note

In case your machine has an unsupported network adapter device, no firmware update will occur and one of the error messages below will be printed. Please contact your hardware vendor for help with firmware updates.

Error message #1:

Device #1:

Device Type: ConnectX4

Part Number: MCX456A-ECA

Description: ConnectX-4 VPI adapter card; EDR IB (100Gb/s) and 100GbE; dual-port QSFP28; PCIe3.0 x16; ROHS R6

PSID: MT_2190110032

PCI Device Name: 0b:00.0

Base MAC: 0000e41d2d5cf810

Versions: Current Available

FW 12.14.0114 N/A

Status: No matching image found

Error message #2:

The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033)

To obtain firmware for this device, please contact your HW vendor.

4. **Case A**: If the installation script has performed a firmware update on your network adapter, you need to either restart the driver or reboot your system before the firmware update can take effect. Refer to the table below to find the appropriate action for your specific card.

Action \ Adapter	Driver Restart	Standard Reboot (Soft Reset)	Cold Reboot (Hard Reset)
Standard ConnectX- 4/ConnectX-4 Lx or higher	-	+	-
Adapters with Multi-Host Support	_	-	+
Socket Direct Cards	-	-	+

Case B: If the installations script has not performed a firmware upgrade on your network adapter, restart the driver by running: "/etc/init.d/openibd restart".

- 5. (InfiniBand only) Run the hca_self_test.ofed utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:
- HCA firmware version
- Kernel architecture
- Driver version
- Number of active HCA ports along with their states
- Node GUID

For more details on hca_self_test.ofed, see the file docs/readme_and_user_manual/hca_self_test.readme.

After installation completion, information about the OFED installation, such as prefix, kernel version, and installation parameters can be retrieved by running the command /etc/infiniband/info. Most of the OFED components can be configured or reconfigured after the installation, by modifying the relevant configuration files. See the relevant chapters in this manual for details.

The list of the modules that will be loaded automatically upon boot can be found in the /etc/infiniband/openib.conf file.

(i) Note

Installing OFED will replace the RDMA stack and remove existing 3rd party RDMA connectors.

Installation Results

Soft ware	 Most of MLNX_OFED packages are installed under the "/usr" directory except for the following packages which are installed under the "/opt" directory: fca and ibutils iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma The kernel modules are installed under /lib/modules/`uname -r`/updates on SLES and Fedora Distributions /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions /lib/modules/`uname -r`/updates/dkms/ on Ubuntu 	
Firm ware	Expansion Norminage, the automatic minimare update will also bu	

Installation Logging

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file.

The path to the directory containing the log files will be displayed after running the installation script in the following format:

Example:

Logs dir: /tmp/MLNX_OFED_LINUX-4.4-1.0.0.0.IBMM2150110033.logs

Driver Load Upon System Boot

Upon system boot, the NVIDIA drivers will be loaded automatically.

To prevent the automatic load of the NVIDIA drivers upon system boot:

≻

1. Add the following lines to the "/etc/modprobe.d/mlnx.conf" file.

blacklist mlx5_core blacklist mlx5_ib

- 2. Set "ONBOOT=no" in the "/etc/infiniband/openib.conf" file.
- 3. If the modules exist in the initramfs file, they can automatically be loaded by the kernel. To prevent this behavior, update the initramfs using the operating systems' standard tools. Note: The process of updating the initramfs will add the blacklists from step 1, and will prevent the kernel from loading the modules automatically.

mInxofedinstall Return Codes

The table below lists the mInxofedinstall script return codes and their meanings.

Return Code	Meaning
0	The Installation ended successfully
1	The installation failed
2	No firmware was found for the adapter device
22	Invalid parameter
28	Not enough free space
171	Not applicable to this system configuration. This can occur when the required hardware is not present on the system
172	Prerequisites are not met. For example, missing the required software installed or the hardware is not configured correctly
173	Failed to start the mst driver

Soft ware	 Most of MLNX_OFED packages are installed under the "/usr" directory except for the following packages which are installed under the "/opt" directory: fca and ibutils iproute2 (rdma tool) - installed under /opt/Mellanox/iproute2/sbin/rdma The kernel modules are installed under /lib/modules/`uname -r`/updates on SLES and Fedora Distributions /lib/modules/`uname -r`/extra/mlnx-ofa_kernel on RHEL and other RedHat like Distributions /lib/modules/`uname -r`/updates/dkms/ on Ubuntu
Firm ware	 The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled: The installation script is run in default mode; that is, without the option 'without- fw-update' The firmware version of the adapter device is older than the firmware version included with the OFED ISO image Note: If an adapter's Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image. In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed. "The firmware for this device is not distributed inside NVIDIA driver: 0000:01:00.0 (PSID: IBM2150110033) To obtain firmware for this device, please contact your HW vendor."

Installation Logging

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file.

The path to the directory containing the log files will be displayed after running the installation script in the following format:

Example:

Logs dir: /tmp/MLNX_OFED_LINUX-4.4-1.0.0.0.IBMM2150110033.logs

Driver Load Upon System Boot

Upon system boot, the NVIDIA drivers will be loaded automatically.

To prevent the automatic load of the NVIDIA drivers upon system boot:

≻

1. Add the following lines to the "/etc/modprobe.d/mlnx.conf" file.

blacklist mlx5_core blacklist mlx5_ib

2. Set "ONBOOT=no" in the "/etc/infiniband/openib.conf" file.

3. If the modules exist in the initramfs file, they can automatically be loaded by the kernel. To prevent this behavior, update the initramfs using the operating systems' standard tools. **Note**: The process of updating the initramfs will add the blacklists from step 1, and will prevent the kernel from loading the modules automatically.

mInxofedinstall Return Codes

The table below lists the mInxofedinstall script return codes and their meanings.

Installation on Community Operating Systems

NVIDIA provides OFED packages to be installed on common operating systems. These packages are provided as binaries, and NVIDIA provided full support for them. This model is now referred to as "Primary support".

Starting OFED 5.6, NVIDIA is introducing a new support model for OFED used on open source community operating systems. The goal of this new support model is to enable customers to use community-maintained variants of the Linux operating system, without being limited to major distributions that NVIDIA provides primary support for.

In the community model, there is shared responsibility between NVIDIA and customers choosing to use community operating systems in their environment. NVIDIA owns basic validation for the operating systems, so that customers know they can expect OFED to work. Customers are responsible for building their own packages and binaries (based on source code and build instructions detailed below), and can also choose to deploy parts of OFED instead of the whole package.

In case of issues, for customers that are entitled for NVIDIA support (e.g. customers who have an applicable support contract), NVIDIA will do the best effort to assist, but may require the customer to work with the community to fix issues that are deemed to be caused by the community breaking OFED, as opposed to NVIDIA owning the fix end to end.

Overall, the following should be noted when running OFED on the community-supported operating systems:

- NVIDIA will perform sanity testing of OFED on community supported OSs
- NVIDIA will declare which kernel versions were tested (min/max), based on mainstream kernel (kernel.org) versions
- NVIDIA will not ship binaries/installation packages for community-supported OSs
- Customers will use source code of OFED and build guidance and will need to build their own binaries & installation packages
- Customers will be able to pick and choose the parts of OFED they deploy (e.g., drivers only, user tools only, and so forth)
- In case of bug reports, customers may be asked to reproduce on a primarysupported distribution (such as RHEL) in order to have support
- If issues are deemed specific to the community OS (e.g., if an OS deviates from mainstream in a way that breaks OFED), it is the customer's responsibility to work with the community to fix it

Below are the instructions how to build OFED from the sources provided by NVIDIA.

Installing OFED on Community Operating Systems

- 1. Download sources from <u>https://network.nvidia.com/products/infiniband-</u> <u>drivers/linux/mlnx_ofed</u>
 - 1. Scroll down to the and click the "Download" tab.
 - 2. Choose the relevant package, depending on the host operating system (the package format is what really matters (rpm/deb)).
 - 3. Download the desired tgz SOURCE package (e.g., MLNX_OFED_SRC-<debian?>-<ver>.tgz)
- 2. Unpack the tarball.

tar -xzf MLNX_OFED_SRC-<debian?>-<version>.tgz

3. Go to the extracted directory.

cd MLNX_OFED_SOURCE-5.6-x.x.x.

4. Run the installation script with the relevant options.

./install.pl <option 1> <option 2> .

Use "./install.pl --help" to see all the options and choose the desired option. For example: "./install.pl --all" should build and install all the default packages.

5. While running install.pl, the script may fail on missing dependencies. Those dependencies should be installed manually before running install.pl again.

Proprietary Packages

The installation procedure does not install proprietary packages—propriety packages should be installed upon request.

List of closed source proprietary packages:

- Clusterkit
- DPCP
- hcoll
- Sharp
- ibutils2
- opensm

Today the only way to install these packages is by using an already-built rpm/deb from a similar primary operating system.

The following table maps the community OSs that are most similar to Primary OSs based on internal testing:

Community OS	Most Similar Primary OS
BCLINUX7.6	RHEL 7.9
BCLINUX7.7	RHEL 7.9
BCLINUX8.1	RHEL 8.5
Debian9.13	Debian 10.8
EulerOS2.0sp5	EulerOS2.0sp10
EulerOS2.0sp8	EulerOS2.0sp10
RHEL/CentOS7.5	RHEL 7.9
RHEL/CentOS7.5alternate	RHEL 7.9
RHEL/CentOS7.6alternate	RHEL 7.9
RHEL/CentOS7.7	RHEL 7.9
RHEL/CentOS7.8	RHEL 7.9

Community OS	Most Similar Primary OS
SLES12SP2	SLES12 SP5
SLES12SP3	SLES12 SP5
SLES12SP4	SLES12 SP5
Ubuntu16.04	Ubuntu22.04
Alma 8.5	RHEL 8.5
Anolis OS 8.4	RHEL 8.5
CentOS Strea m 8.5	RHEL 8.5
Fedora 35	RHEL 8.5
OpenEuler OS 22.03 LTS	OpenEuler20 SP3
OpenSUSE 15.3	SLES15 SP3
Photon OS 3.0	RHEL 7.9
Rocky 8.5	RHEL 8.5
Rocky 8.6	RHEL 8.6

Download and install MFT and Firmware

NVIDIA Firmware Tools (MFT) Download

<u>https://network.nvidia.com/products/adapter-software/firmware-tools</u>

Firmware Downloads

• <u>https://network.nvidia.com/support/firmware/firmware-downloads</u>

Creating Metapackages and Installing OFED on Multiple Servers (RPM-Based Systems)

By building the packages using install.pl, the RPMs will be created under /tmp/OFED-internal-5.6x.x.x/RPMS/<OS>/<arch>/

1. List the packages that will be installed by running the following:

```
/tmp/OFED-internal-5.6-x.x.x/install.pl--all -p
```

2. Create working directories.

Example:

```
mkdir -p /tmp/OFED_topdir/SOURCES/<rpmname>-<version>
mkdir -p /tmp/OFED_topdir/SPECS
mkdir -p /tmp/OFED_topdir/BUILD
mkdir -p /tmp/OFED_topdir/SRPMS
mkdir -p /tmp/OFED_topdir/RPMS
```

For --all option it used to set *rpmname* as mlnx-ofed-all

OFED will be set to the current GA version being used.

3. Create a new XML file under /tmp/OFED_topdir/comp.xml that contains all the required packages.

Example:

```
<?xml version="1.0" encoding="UTF-8"?>
<comps>
<group>
<id>all</id>
<name>MLNX_OFED ALL</name>
<default>true</default>
<description>Mellanox OpenFabrics Enterprise Distribution for
Linux: MLNX_OFED ALL packages </description>
<uservisible>true</uservisible>
<packagelist>
```

```
/* ----- The list below needs to be aligned with the packages listed by "install.pl --all -p" ----- */
<packagereq type="default">mlnx-tools</packagereq>
<packagereq type="default">rdma-core-devel</packagereq>
<packagereq type="default">perftest</packagereq>
<packagereq type="default">Package Example 1</packagereq>
<packagereq type="default">Package Example 2</packagereq>
<packagereq type="default">Package Example 2</packagereq>
<packagereq type="default">mlnx-ofed-all</packagereq>
<packagereq type="default">Package Example 2</packagereq>
</packagereq>
</packagereq type="default">Package Example 2</packagereq>
```

4. Create tarball with the correct source name:

```
echo <rpmname>-<version> >
/tmp/OFED_topdir/SOURCES/<rpmname>-<version>/<rpmname>-
release
cd /tmp/OFED_topdir/SOURCES
tar czf <rpmname>-<version>.tar.gz <rpmname>-<version>
```

- 5. Create new spec file that requires all the packages in that directory under /tmp/OFED_topdir/SPECS/<rpmname>.spec
- 6. Build the RPM metapackage from the spec file that was created before. As a result, the RPMs will be created under /tmp/OFED_topdir/RPMS/<arch>/

```
rpmbuild -ba --target noarch --define '_source_filedigest_algorithm md5' -
-define '_binary_filedigest_algorithm md5' --define '_topdir /tmp/OFED_topdir' --
define '_sourcedir %{_topdir}/SOURCES' --define '_specdir %{_topdir}/SPECS' --
```

```
define '_srcrpmdir %{_topdir}/SRPMS' --define '_rpmdir %{_topdir}/RPMS'
/tmp/OFED_topdir/SPECS/<rpmname>.spec
```

7. Copy the newly-created RPMs into the RPMS directory that was created before.

```
cp /tmp/OFED_topdir/RPMS/noarch/mlnx-ofed-all-
5.6-.rhel8u3.noarch.rpm /tmp/OFED-internal-5.6-
x.x.x/<arch>/redhat-release-8.3-1.0.el8/x86_64/
```

8. Create repodata.

```
createrepo -q -g /tmp/OFED_topdir/comps.xml /tmp/OFED-
internal-5.6-x.x.x/RPMS/redhat-release-8.3-1.0.el8/<arch>/
```

- 9. Copy the generated RPMS to any server and do the following to install OFED:
 - 1. Create new yum repository.

```
cat /etc/yum.repos.d/mlnx_ofed.repo
[mlnx_ofed]
name=mlnx_ofed
baseurl=/tmp/OFED-internal-5.6-x.x.x/RPMS/OS/<arch>/
enabled=1
gpgcheck=0
```

2. Refresh repository list.

yum repolist

3. Install OFED.

yum install mlnx-ofed-all

Creating Metapackages and Installing OFED on Multiple Servers (Deb-Based Systems)

By building the packages using install.pl, the debs will be created under /tmp/OFED-internal-5.6-x.x.x/DEBS/<OS>/<arch>/

1. List the packages that will be installed by the following:

/tmp/OFED-internal-debian-5.6-x.x.x/install.pl --all -p

2. Create working directories.

Example:

```
mkdir -p /tmp/OFED_topdir/<rpmname>-<version>/debian/source
echo <rpmname>-<version> > /tmp/OFED_topdir/<rpmname>-
<version>/<rpmname>-release
```

For --all option it used to set *rpmname* as mlnx-ofed-all

For OFED 5.6, the version will be set to 5.6

3. Create rules file under /tmp/OFED_topdir/<rpmname>-<version>/debian/rules

```
#!/usr/bin/make -f
# -*- makefile -*-
```

4. Change rules file mode.

chmod 755 /tmp/OFED_topdir/<rpmname>-<version>/debian/rules

5. Create rules file under /tmp/OFED_topdir/<rpmname>-<version>/debian/compat

echo 9 > /tmp/OFED_topdir/<rpmname>-<version>/debian/compat

6. Create source/format file under /tmp/OFED_topdir/<rpmname>-<version>/debian/source/format

echo "3.0(quilt)" > /tmp/OFED_topdir/<rpmname><version>/debian/source/format

7. Create changelog file under /tmp/OFED_topdir/<rpmname>-<version>/debian/changelog

Exmaple:

mlnx-ofed-all (5.6-x.x.x) unstable; urgency=low

```
* Initial release
-- your username <mail> Sun, 15 May 2022 21:00:00 +0200
```

8. Create postinst script (if needed) under /tmp/OFED_topdir/<rpmname>-<version>/debian/<rpmname>.postinst

```
#!/bin/bash
cd /lib/modules
for dd in `/bin/ls`
do
/sbin/depmod $dd >/dev/null 2>&1
done
if [ -f /usr/bin/ofed_info ]; then
sed -i -r -e "s/^(OFED)(.*)(-[0-9]*.*-[0-9]*.*):/MLNX_OFED_LINUX-5.6-x.x.x.x (\1\3):\n/"
/usr/bin/ofed_info
sed -i -r -e "s/(.*then echo) (.*):(.*)/\1 MLNX_OFED_LINUX-5.6-x.x.x.x: \3/"
/usr/bin/ofed info
sed -i -r -e "s/(.*X-n\"]; then echo) (.*)(; exit.*)/\1 5.6-x.x.x.x \3/"
/usr/bin/ofed_info
sed -i -e "s/OFED-internal/MLNX_OFED_LINUX/g" /usr/bin/ofed_info
fi
# Switch off opensmd service
/sbin/chkconfig --set opensmd off > /dev/null 2>&1 || true
/sbin/chkconfig opensmd off > /dev/null 2>&1 || true
if [ -f "/etc/init.d/opensmd" ] ; then
if [ -e /sbin/chkconfig ]; then
/sbin/chkconfig --del opensmd > /dev/null 2>&1 || true
elif [ -e /usr/sbin/update-rc.d ]; then
/usr/sbin/update-rc.d -f opensmd remove > /dev/null 2>&1 || true
else
/usr/lib/lsb/remove_initd /etc/init.d/opensmd > /dev/null 2>&1
true
```

```
fi
fi
fi
# Disable ibacm daemon by default
chkconfig --del ibacm > /dev/null 2>&1 || true
# disable SDP and QIB loading by default
if [ -e /etc/infiniband/openib.conf ]; then
sed -i -r -e "s/^SDP_LOAD=.*/SDP_LOAD=no/"
/etc/infiniband/openib.conf
sed -i -r -e "s/^QIB_LOAD=.*/QIB_LOAD=no/"
/etc/infiniband/openib.conf
fi
/sbin/ldconfig > /dev/null 2>&1 || true
```

9. Create control file under /tmp/OFED_topdir/<rpmname>-<version>/debian/control

```
Source: mlnx-ofed-all
Section: utils
Priority: extra
Maintainer: your username <mail>
Build-Depends: debhelper (>= 9.0.0)
Standards-Version: 3.9.2
Homepage: http://www.mellanox.com
Package:mlnx-ofed-all // PACKAGE NAME
Architecture: all
```

// list of the packages we need to install in this group (see the list of packages to install we created at first)

```
Depends: ${shlibs:Depends}, ${misc:Depends}, dpcp (>=1.1.25-
1.56076), mstflint (>=4.16.0-1.56076), rdmacm-utils (>=56mlnx40-
```

```
1.56076), mlnx-tools (>=5.2.0-0.56076), mlnx-iproute2 (>=5.16.0-
1.56076), opensm (>=5.11.0.MLNX20220418.fd3d650-0.1.56076),
ibutils2 (>=2.1.1-0.148.MLNX20220418.g60b8156.56076), ofed-
scripts (>=5.6-OFED.5.6.0.7.6), dump-pr (>=1.0-
5.11.0.MLNX20220418.g7e9d922.56076), mlnx-ethtool (>=5.15-1.56076),
perftest (>=4.5-0.14.gd962d8c.56076), libibverbs-dev
(>=56mlnx40-1.56076), ucx (>=1.13.0-1.56076), ibsim-doc (>=0.10-
1.56076), srp-dkms (>=5.6-OFED.5.6.0.7.6.1), ibacm (>=56mlnx40-
1.56076), sharp (>=2.7.0.MLNX20220331.8d57397a-1.56076), libibmad5
(>=56mlnx40-1.56076), mlnx-ofed-kernel-utils (>=5.6-
OFED.5.6.0.7.6.1), srptools (>=56mlnx40-1.56076), ibsim (>=0.10-
1.56076), libibmad-dev (>=56mlnx40-1.56076), libibumad3
(>=56mlnx40-1.56076), ibverbs-utils (>=56mlnx40-1.56076), rdma-
core (>=56mlnx40-1.56076), libibverbs1 (>=56mlnx40-1.56076),
isert-dkms (>=5.6-OFED.5.6.0.7.6.1), libibumad-dev (>=56mlnx40-
1.56076), iser-dkms (>=5.6-OFED.5.6.0.7.6.1), ibdump (>=6.0.0-
1.56076), libopensm-devel (>=5.11.0.MLNX20220418.fd3d650-
0.1.56076), libopensm (>=5.11.0.MLNX20220418.fd3d650-0.1.56076),
kernel-mft-dkms (>=4.20.0-29), libibnetdisc5 (>=56mlnx40-
1.56076), librdmacm1 (>=56mlnx40-1.56076), opensm-doc
(>=5.11.0.MLNX20220418.fd3d650-0.1.56076), mlnx-ofed-kernel-dkms
(>=5.6-OFED.5.6.0.7.6.1), infiniband-diags (>=56mlnx40-1.56076),
mft (>=4.15.1-9), librdmacm-dev (>=56mlnx40-1.56076), ibverbs-
providers (>=56mlnx40-1.56076)
Description: MLNX_OFED all installer package (with DKMS
support)
Package:mlnx-ofed-all-exact
Architecture: all
```

10. Create copyright file under /tmp/OFED_topdir/<rpmname>-<version>/debian/copyright

On Debian systems, a copy of the General Public License version 2 could be found under /usr/share/common-licenses/GPL-2

Format: https://www.debian.org/doc/packaging-manuals/copyright-format/1.0 Files: * Copyright: 2022, NVIDIA Corporation License: GPLv2-and-2BSD * This software is available to you under a choice of one of two * licenses. You may choose to be licensed under the terms of the GNU * General Public License (GPL) Version 2, available from the file * COPYING in the main directory of this source tree, or the * BSD license below: \star * Redistribution and use in source and binary forms, with or * without modification, are permitted provided that the following * conditions are met: * * - Redistributions of source code must retain the above * copyright notice, this list of conditions and the following * disclaimer. * * - Redistributions in binary form must reproduce the above * copyright notice, this list of conditions and the following * disclaimer in the documentation and/or other materials * provided with the distribution. \star * THE SOFTWARE IS PROVIDED "ASIS", WITHOUT WARRANTY OF ANY KIND, * EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF * MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND * NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS

```
* BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY,
WHETHER IN AN
* ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF
OR IN
* CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS
IN THE
* SOFTWARE.
```

11. Creating tarball.

```
cd /tmp/OFED_topdir
tar czf <rpmname>_<version>.orig.tar.gz <rpmname>-<version>
cd <rpmname>-<version>
dpkg-buildpackage -us -uc
cp -af *.deb /tmp/OFED-internal-5.6-x.x.x/DEBS/OS/arch/
```

12. Create new yum repository under /etc/yum.repos.d/mlnx_ofed.repo

```
[mlnx_ofed]
deb [trusted=yes] file:/tmp/OFED-internal-5.6-
0.7.6/DEBS/debian10.8/x86_64/ ./
```

13. Refresh repository list.

apt update

14. Install OFED.

apt install mlnx-ofed-all

i) Note

mpitests package is disabled on Fedora 35 and Photon due to building issues.mpitests package is disabled on Fedora 35 and Photon due to building issues.

Additional Installation Procedures

Installing MLNX_OFED Using YUM

This type of installation is applicable to RedHat/OL and Fedora operating systems.

Setting up MLNX_OFED YUM Repository

- 1. Log into the installation machine as root.
- 2. Mount the ISO image on your machine and copy its content to a shared location in your network.

mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU
arch>.iso /mnt

3. Download and install NVIDIA's GPG-KEY:

The key can be downloaded via the following link:

http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox-SHA256

```
# wget http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox-SHA256
--2018-01-25 13:52:30-- http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-
Mellanox-SHA256
Resolving www.mellanox.com.. 72.3.194.0
Connecting to www.mellanox.com[72.3.194.0]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1354 (1.3K) [text/plain]
Saving to: ?RPM-GPG-KEY-Mellanox-SHA256?
100%[======>] 1,354
--.-K/s in 0s
2018-01-25 13:52:30 (247 MB/s) - ?RPM-GPG-KEY-Mellanox-SHA256?
saved [1354/1354]
```

4. Install the key.

```
# sudo rpm --import RPM-GPG-KEY-Mellanox-SHA256
warning: rpmts_HdrFromFdno: Header V3 DSA/SHA1 Signature, key
ID 6224c050: NOKEY
Retrieving key from file:///repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-
Mellanox-SHA256
Importing GPG key 0x6224C050:
   Userid: "Mellanox Technologies (Mellanox Technologies - Signing Key v2)
<support@mellanox.com>"
   From : /repos/MLNX_OFED/<MLNX_OFED file>/RPM-GPG-KEY-
Mellanox-SHA256
Is this ok [y/N]:
```

5. Check that the key was successfully imported.

```
# rpm -q gpg-pubkey --qf '%{NAME}-%{VERSION}-%{RELEASE}\t%{SUMMARY}\n' |
grep Mellanox
gpg-pubkey-a9e4b643-520791ba gpg(Mellanox Technologies
<support@mellanox.com>)
```

6. Create a yum repository configuration file called "/etc/yum.repos.d/mlnx_ofed.repo" with the following content:

```
[mlnx_ofed]
name=MLNX_OFED Repository
baseurl=file:///<path to extracted MLNX_OFED package>/RPMS
enabled=1
gpgkey=file:///<path to the downloaded key RPM-GPG-KEY-Mellanox-SHA256>
gpgcheck=1
```

7. Check that the repository was successfully added.

```
# yum repolist
Loaded plugins: product-id, security, subscription-manager
This system is not registered to Red Hat Subscription
Management. You can use subscription-manager to register.
repo id repo name
status
mlnx_ofed MLNX_OFED Repository
108
rpmforge RHEL 6Server - RPMforge.net - dag
4,597
repolist: 8,351
```

Setting up MLNX_OFED YUM Repository Using --add-kernel-support

- 1. Log into the installation machine as root.
- 2. Mount the ISO image on your machine and copy its content to a shared location in your network.

```
# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU
arch>.iso /mnt
```

3. Build the packages with kernel support and create the tarball.

```
# /mnt/mlnx_add_kernel_support.sh --make-tgz <optional --kmp>
-k $(uname -r) -m /mnt/
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.6
under /tmp directory.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_iso.4120_logs/mlnx_ofed_iso.4120.log
Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMS . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a
repository, then please note
WARNING: that it contains unsigned rpms, therefore, you need
to disable the gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

cd /tmp/

```
# tar -xvf /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-
ext.tgz
```

5. Create a YUM repository configuration file called "/etc/yum.repos.d/mlnx_ofed.repo" with the following content:

```
[mlnx_ofed]
name=MLNX_OFED Repository
baseurl=file:///<path to extracted MLNX_OFED package>/RPMS
enabled=1
gpgcheck=0
```

6. Check that the repository was successfully added.

```
# yum repolist
Loaded plugins: product-id, security, subscription-manager
This system is not registered to Red Hat Subscription
Management. You can use subscription-manager to register.
repo id repo name
status
mlnx_ofed MLNX_OFED Repository
108
rpmforge RHEL 6Server - RPMforge.net - dag
4,597
repolist: 8,351
```

Installing MLNX_OFED Using the YUM Tool

After setting up the YUM repository for MLNX_OFED package, perform the following:

1. View the available package groups by invoking:

```
# yum search mlnx-ofed-
mlnx-ofed-all.noarch : MLNX_OFED all installer package (with
KMP support)
mlnx-ofed-all-user-only.noarch : MLNX_OFED all-user-only
installer package (User Space packages only)
mlnx-ofed-basic.noarch : MLNX_OFED basic installer package
(with KMP support)
mlnx-ofed-basic-user-only.noarch : MLNX_OFED basic-user-only
installer package (User Space packages only)
mlnx-ofed-bluefield.noarch : MLNX_OFED bluefield installer
package (with KMP support)
mlnx-ofed-bluefield-user-only.noarch : MLNX_OFED bluefield-
user-only installer package (User Space packages only)
mlnx-ofed-dpdk.noarch : MLNX_OFED dpdk installer package (with
KMP support)
mlnx-ofed-dpdk-upstream-libs.noarch : MLNX_OFED dpdk-
upstream-libs installer package (with KMP support)
mlnx-ofed-dpdk-upstream-libs-user-only.noarch : MLNX_OFED
dpdk-upstream-libs-user-only installer package (User Space
packages only)
mlnx-ofed-dpdk-user-only.noarch : MLNX_OFED dpdk-user-only
installer package (User Space packages only)
mlnx-ofed-eth-only-user-only.noarch : MLNX_OFED eth-only-
user-only installer package (User Space packages only)
mlnx-ofed-guest.noarch : MLNX_OFED guest installer package
(with KMP support)
mlnx-ofed-guest-user-only.noarch : MLNX_OFED guest-user-only
installer package (User Space packages only)
mlnx-ofed-hpc.noarch : MLNX_OFED hpc installer package (with
KMP support)
mlnx-ofed-hpc-user-only.noarch : MLNX_OFED hpc-user-only
installer package (User Space packages only)
```

```
mlnx-ofed-hypervisor.noarch : MLNX_OFED hypervisor installer
package (with KMP support)
mlnx-ofed-hypervisor-user-only.noarch : MLNX_OFED hypervisor-
user-only installer package (User Space packages only)
mlnx-ofed-kernel-only.noarch : MLNX_OFED kernel-only
installer package (with KMP support)
mlnx-ofed-vma.noarch : MLNX_OFED vma installer package (with
KMP support)
mlnx-ofed-vma-eth.noarch : MLNX_OFED vma-eth installer package
(with KMP support)
mlnx-ofed-vma-eth-user-only.noarch : MLNX_OFED vma-eth-user-
only installer package (User Space packages only)
mlnx-ofed-vma-user-only.noarch : MLNX_OFED vma-user-only
installer package (User Space packages only)
mlnx-ofed-vma-vpi.noarch : MLNX_OFED vma-vpi installer package
(with KMP support)
mlnx-ofed-vma-vpi-user-only.noarch : MLNX_OFED vma-vpi-user-
only installer package (User Space packages only
```

where:

mlnx-ofed-all	Installs all available packages in MLNX_OFED	
mlnx-ofed-basic	Installs basic packages required for running NVIDIA cards	
mlnx-ofed-guest	Installs packages required by guest OS	
mlnx-ofed-hpc	Installs packages required for HPC	
mlnx-ofed-hypervisor	Installs packages required by hypervisor OS	
mlnx-ofed-vma	Installs packages required by VMA	
mlnx-ofed-vma-eth	Installs packages required by VMA to work over Ethernet	
mlnx-ofed-vma-vpi	Installs packages required by VMA to support VPI	
bluefield	Installs packages required for BlueField	
dpdk	Installs packages required for DPDK	
dpdk-upstream-libs	Installs packages required for DPDK using RDMA-Core	

kernel-only	Installs packages required for a non-default kernel
-------------	---

Note: MLNX_OFED provides kernel module RPM packages with KMP support for RHEL and SLES. For other operating systems, kernel module RPM packages are provided only for the operating system's default kernel. In this case, the group RPM packages have the supported kernel version in their package's name.

Example:

mlnx-ofed-all-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED all installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP) support) mlnx-ofed-basic-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED basic installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support) mlnx-ofed-guest-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED guest installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support) mlnx-ofed-hpc-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hpc installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP) support) mlnx-ofed-hypervisor-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hypervisor installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support) mlnx-ofed-vma-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP) support) mlnx-ofed-vma-eth-3.17.4-301.fc21.x86 64.noarch : MLNX OFED vma-eth installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support) mlnx-ofed-vma-vpi-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma-vpi installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support) mlnx-ofed-hypervisor-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED hypervisor installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP support)

```
mlnx-ofed-vma-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED vma
installer package for kernel 3.17.4-301.fc21.x86_64 (without KMP
support)
mlnx-ofed-vma-eth-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED
vma-eth installer package for kernel 3.17.4-301.fc21.x86_64
(without KMP support)
mlnx-ofed-vma-vpi-3.17.4-301.fc21.x86_64.noarch : MLNX_OFED
vma-vpi installer package for kernel 3.17.4-301.fc21.x86_64
(without KMP support)
```

When using an operating system different than RHEL or SLES, or you have installed a kernel that is not supported by default in MLNX_OFED, you can use the mlnx_add_kernel_support.sh script to build MLNX_OFED for your kernel.

The script will automatically build the matching group RPM packages for your kernel so that you can still install MLNX_OFED via yum.

Please note that the resulting MLNX_OFED repository will contain unsigned RPMs, therefore, you should set 'gpgcheck=0' in the repository configuration file.

2. Install the desired group.

```
# yum install mlnx-ofed-all
Loaded plugins: langpacks, product-id, subscription-manager
Resolving Dependencies
--> Running transaction check
--> Package mlnx-ofed-all.noarch 0:3.1-0.1.2 will be installed
--> Processing Dependency: kmod-isert = 1.0-
OFED.3.1.0.1.2.1.g832a737.rhel7u1 for package: mlnx-ofed-all-3.1-
0.1.2.noarch
.....
qperf.x86_64 0:0.4.9-9
rds-devel.x86_64 0:2.0.7-1.12
rds-tools.x86_64 0:2.0.7-1.12
sdpnetstat.x86_64 0:1.60-26
```

```
srptools.x86_64 0:1.0.2-12
```

Complete!

i) Note

Installing MLNX_OFED using the "YUM" tool does not automatically update the firmware.

To update the firmware to the version included in MLNX_OFED package, run:

yum install mlnx-fw-updater

OR:

Update the firmware to the latest version available on NVIDIA's website as described in "<u>Updating Firmware After Installation</u>" section.

Installing MLNX_OFED Using apt-get

This type of installation is applicable to Debian and Ubuntu operating systems.

Setting up MLNX_OFED apt-get Repository

- 1. Log into the installation machine as root.
- 2. Extract the MLNX_OFED package on a shared location in your network.

It can be downloaded from <u>https://www.nvidia.com/en-us/networking/</u> \rightarrow Products \rightarrow Software \rightarrow InfiniBand Drivers.

3. Create an apt-get repository configuration file called "/etc/apt/sources.list.d/mlnx_ofed.list" with the following content: deb file:/<path to extracted MLNX_OFED package>/DEBS ./

4. Download and install NVIDIA's Technologies GPG-KEY.

```
# wget -q0 - http://www.mellanox.com/downloads/ofed/RPM-GPG-KEY-Mellanox-
SHA256-SHA256 | sudo apt-key add -
```

5. Verify that the key was successfully imported.

```
# apt-key list
pub 1024D/A9E4B643 2013-08-11
uid Mellanox Technologies <support@mellanox.com>
sub 1024g/09FCC269 2013-08-11
```

6. Update the apt-get cache.

```
# sudo apt-get update
```

Setting up MLNX_OFED apt-get Repository Using --add-kernel-support

- 1. Log into the installation machine as root.
- 2. Mount the ISO image on your machine and copy its content to a shared location in your network.

```
# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU
arch>.iso /mnt
```

3. Build the packages with kernel support and create the tarball.

```
# /mnt/mlnx_add_kernel_support.sh --make-tgz <optional --kmp>
-k $(uname -r) -m /mnt/
Note: This program will create MLNX_OFED_LINUX TGZ for rhel7.6
under /tmp directory.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_iso.4120_logs/mlnx_ofed_iso.4120.log
Checking if all needed packages are installed...
Building MLNX_OFED_LINUX RPMS . Please wait...
Creating metadata-rpms for 3.10.0-957.21.3.el7.x86_64 ...
WARNING: If you are going to configure this package as a
repository, then please note
WARNING: that it contains unsigned rpms, therefore, you need
to disable the gpgcheck
WARNING: by setting 'gpgcheck=0' in the repository conf file.
Created /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-ext.tgz
```

4. Open the tarball.

```
# cd /tmp/
# tar -xvf /tmp/MLNX_OFED_LINUX-5.2-0.5.5.0-rhel7.6-x86_64-
ext.tgz
```

5. Create an apt-get repository configuration file called "/etc/apt/sources.list.d/mlnx_ofed.list" with the following content:

```
deb [trusted=yes] file:/<path to extracted MLNX_OFED
package>/DEBS ./
```

6. Update the apt-get cache.

Installing MLNX_OFED Using the apt-get Tool

After setting up the apt-get repository for MLNX_OFED package, perform the following:

1. View the available package groups by invoking:

```
# apt-cache search mlnx-ofed-
apt-cache search mlnx-ofed .....
knem-dkms - DKMS support for mlnx-ofed kernel modules
mlnx-ofed-kernel-dkms - DKMS support for mlnx-ofed kernel
modules
mlnx-ofed-kernel-utils - Userspace tools to restart and tune
mlnx-ofed kernel modules
mlnx-ofed-vma-vpi - MLNX_OFED vma-vpi installer package (with
DKMS support)
mlnx-ofed-kernel-only - MLNX_OFED kernel-only installer package
(with DKMS support)
mlnx-ofed-bluefield - MLNX_OFED bluefield installer package
(with DKMS support)
mlnx-ofed-hpc-user-only - MLNX_OFED hpc-user-only installer
package (User Space packages only)
mlnx-ofed-dpdk-user-only - MLNX_OFED dpdk-user-only installer
package (User Space packages only)
mlnx-ofed-all-exact - MLNX_OFED all installer package
                                                      (with
DKMS support) (exact)
mlnx-ofed-all - MLNX_OFED all installer package (with DKMS
support)
mlnx-ofed-vma-vpi-user-only - MLNX_OFED vma-vpi-user-only
installer package (User Space packages only)
```

```
mlnx-ofed-eth-only-user-only - MLNX_OFED eth-only-user-only
installer package (User Space packages only)
mlnx-ofed-vma-user-only - MLNX_OFED vma-user-only installer
package (User Space packages only)
mlnx-ofed-hpc - MLNX_OFED hpc installer package (with DKMS)
support)
mlnx-ofed-bluefield-user-only - MLNX_OFED bluefield-user-only
installer package (User Space packages only)
mlnx-ofed-dpdk - MLNX_OFED dpdk installer package (with DKMS
support)
mlnx-ofed-vma-eth-user-only - MLNX_OFED vma-eth-user-only
installer package (User Space packages only)
mlnx-ofed-all-user-only - MLNX_OFED all-user-only installer
package (User Space packages only)
mlnx-ofed-vma-eth - MLNX_OFED vma-eth installer package
                                                        (with
DKMS support)
mlnx-ofed-vma - MLNX_OFED vma installer package (with DKMS)
support)
mlnx-ofed-dpdk-upstream-libs-user-only - MLNX_OFED dpdk-
upstream-libs-user-only installer package (User Space packages
only)
mlnx-ofed-basic-user-only - MLNX_OFED basic-user-only
installer package (User Space packages only)
mlnx-ofed-basic-exact - MLNX_OFED basic installer package
(with DKMS support) (exact)
mlnx-ofed-basic - MLNX_OFED basic installer package (with DKMS)
support)
mlnx-ofed-dpdk-upstream-libs - MLNX_OFED dpdk-upstream-libs
installer package (with DKMS support)
```

where:

mlnx-ofed-all	MLNX_OFED all installer package
mlnx-ofed-basic	MLNX_OFED basic installer package

mlnx-ofed-vma	MLNX_OFED vma installer package			
mlnx-ofed-hpc	MLNX_OFED HPC installer package			
mlnx-ofed-vma-eth	MLNX_OFED vma-eth installer package			
mlnx-ofed-vma-vpi	MLNX_OFED vma-vpi installer package			
knem-dkms	MLNX_OFED DKMS support for mlnx-ofed kernel modules			
kernel-dkms	MLNX_OFED kernel-dkms installer package			
kernel-only	MLNX_OFED kernel-only installer package			
bluefield	MLNX_OFED bluefield installer package			
mlnx-ofed-all-exact	MLNX_OFED mlnx-ofed-all-exact installer package			
dpdk	MLNX_OFED dpdk installer package			
mlnx-ofed-basic-exact	MLNX_OFED mlnx-ofed-basic-exact installer package			
dpdk-upstream-libs	MLNX_OFED dpdk-upstream-libs installer package			
h				

2. Install the desired group.

apt-get install '<group name>'

Example:

```
apt-get install mlnx-ofed-all
```

(i) Note

Installing MLNX_OFED using the "apt-get" tool does not automatically update the firmware.

To update the firmware to the version included in MLNX_OFED package, run:

apt-get install mlnx-fw-updater

OR:

Update the firmware to the latest version available on NVIDIA's website as described in "<u>Updating Firmware After Installation</u>" section.

Uninstalling the Driver

Uninstalling NVIDIA OFED Using the YUM and apt-get Tools

Use the script /usr/sbin/ofed_uninstall.sh to uninstall the NVIDIA OFED package. The script is part of the ofed-scripts RPM.

Updating Firmware After Installation

The firmware can be updated using one of the following methods.

Updating the Device Online

To update the device online on the machine from the NVIDIA site, use the following command line:

```
mlxfwmanager --online -u -d <device>
```

Example:

```
# mlxfwmanager --online -u -d 0000:01:00.0
Querying Mellanox devices firmware ...
```

Device #1:

Device Type: Part Number: Description: and 200GbE; dual-po		/PI adapter	card; HDR IB (200Gb/s) tall bracket; ROHS R6		
PSID:	MT_000000225				
PCI Device Name:	0000:01:00.0				
Base MAC:	98039b970cc2				
Versions:	Current	Availab	le		
FW	20.26.4012	20.27.1016			
PXE	3.6.0101	3.5.0903			
UEFI	14.21.0016	14.20.0025			
Status:	Up to date				

Updating Firmware and FPGA Image on Innova IPsec Cards

The firmware and FPGA update package (mlnx-fw-updater) are installed under "/opt/mellanox/mlnx-fw-updater" folder.

The latest FW and FPGA update package can be downloaded from <u>nvidia.com/en-</u> <u>us/networking/</u> Products \rightarrow Adapters \rightarrow Smart Adapters \rightarrow Innova IPsec \rightarrow Download tab.



The current update package available on <u>nvidia.com/en-</u> <u>us/networking/</u> does not support the script below. An update package that supports this script will become available in a future release.

You can run the following update script using one of the modes below:

/opt/mellanox/mlnx-fw-updater/mlnx_fpga_updater.sh

• With -u flag to provide URL to the software package (tarball). Example:

./mlnx_fpga_updater.sh -u http://www.mellanox.com/downloads/fpga/ipsec/Innova_IPsec_<version>.tgz

• With -t flag to provide the path to the downloaded tarball. Example:

./mlnx_fpga_updater.sh -t <Innova_IPsec_bundle_file.tgz>

• With -p flag to provide the path to the downloaded and extracted tarball. Example:

./mlnx_fpga_updater.sh -p
<Innova_IPsec_extracted_bundle_directory>

For more information on the script usage, you can run mlnx_fpga_updater.sh -h.



It is recommended to perform firmware and FPGA upgrade on Innova IPsec cards using this script only.

Updating the Device Manually

When running the mInxofedinstall script with the '--without-fw-update'option or using an OEM card that you now wish to (manually) update firmware on your adapter card(s),

perform the steps below. The following steps are also appropriate to burn newer firmware that was downloaded from the website (i.e., <u>nvidia.com/en-us/networking/</u> \rightarrow Support \rightarrow Support \rightarrow Firmware Download).

1. Get the device's PSID.

mlxfwmanager_pci | grep PSID
PSID: MT_1210110019

- 2. Download the firmware BIN file from the website or the OEM website.
- 3. Burn the firmware.

mlxfwmanager_pci -i <fw_file.bin>

4. Reboot the device once the firmware burning is completed.

Updating the Device Firmware Automatically Upon System Boot

Firmware can be automatically updated upon system boot.

The firmware update package (mlnx-fw-updater) is installed in the "/opt/mellanox/mlnxfw-updater" folder, and the openibd service script can invoke the firmware update process if requested on boot.

If the firmware is updated, the following message will be printed to the system's standard logging file:

fw_updater: Firmware was updated. Please reboot your system for the changes to take effect.

Otherwise, the following message will be printed:

fw_updater: Didn't detect new devices with old firmware.

Please note that this feature is disabled by default. To enable the automatic firmware update upon system boot, set the following parameter to "yes" "RUN_FW_UPDATER_ONBOOT=yes" in the openibd service configuration file "/etc/infiniband/openib.conf".

You can opt to exclude a list of devices from the automatic firmware update procedure. To do so, edit the configurations file "/opt/mellanox/mlnx-fw-updater/mlnx-fw-updater.conf" and provide a comma separated list of PCI devices to exclude from the firmware update.

Example:

MLNX_EXCLUDE_DEVICES="00:05.0,00:07.0"

UEFI Secure Boot

All kernel modules included in MLNX_OFED for RHEL and SLES are signed with x.509 key to support loading the modules when Secure Boot is enabled.

Enrolling NVIDIA's x.509 Public Key On your Systems

In order to support loading MLNX_OFED drivers when an OS supporting Secure Boot boots on a UEFI-based system with Secure Boot enabled, the NVIDIA x.509 public key should be added to the UEFI Secure Boot key database and loaded onto the system key ring by the kernel.

Follow these steps below to add the NVIDIA's x.509 public key to your system:



Prior to adding the NVIDIA's x.509 public key to your system, please make sure:

• The 'mokutil' package is installed on your system

- The system is booted in UEFI mode
- 1. Download the x.509 public key.

wget http://www.mellanox.com/downloads/ofed/mlnx_signing_key_pub.der



As of version 23.04, the builds for SLES15 sp4 and sp5 are being singed with a newer signing key. The corresponding public key can be downloaded from <u>https://www.mellanox.com/downloads/ofed/nv_nbu_kernel_signing_key_pub.de</u>

2. Add the public key to the MOK list using the mokutil utility.

You will be asked to enter and confirm a password for this MOK enrollment request.

mokutil --import mlnx_signing_key_pub.der

3. Reboot the system.

The pending MOK key enrollment request will be noticed by shim.efi and it will launch MokManager.efi to allow you to complete the enrollment from the UEFI console. You will need to enter the password you previously associated with this request and confirm the enrollment. Once done, the public key is added to the MOK list, which is persistent. Once a key is in the MOK list, it will be automatically propagated to the system key ring and subsequent will be booted when the UEFI Secure Boot is enabled.



To see what keys have been added to the system key ring on the current boot, install the 'keyutils' package and run: #keyctl list %:.system_keyring

Removing Signature from Kernel Modules

The signature can be removed from a signed kernel module using the 'strip' utility which is provided by the 'binutils' package.

strip -g my_module.ko

The strip utility will change the given file without saving a backup. The operation can be undone only by resigning the kernel module. Hence, we recommend backing up a copy prior to removing the signature.

1. Remove the signature.

```
# rpm -qa | grep -E "kernel-ib|mlnx-ofa_kernel|iser|srp|knem|mlnx-rds|mlnx-
nfsrdma|mlnx-nvme|mlnx-rdma-rxe" | xargs rpm -ql | grep "\.ko$" | xargs
strip -g
```

Once the signature is removed, a message as the below will no longer be presented upon module loading:

"Request for unknown module key 'Mellanox Technologies signing key: 61feb074fc7292f958419386ffdd9d5ca999e403' err -11"

However, please note that a message similar to the following will be presented:

"my_module: module verification failed: signature and/or required key missing - tainting kernel"

This message is presented once, only upon first module boot that either has no signature or whose key is not in the kernel key ring. Therefore, this message may go unnoticed. Once the system is rebooted after unloading and reloading a kernel module, the message will appear (this message cannot be eliminated).

2. Update the initramfs on RHEL systems with the stripped modules.

mkinitrd /boot/initramfs-\$(uname -r).img \$(uname -r) --force

Performance Tuning

Depending on the application of the user's system, it may be necessary to modify the default configuration of network adapters based on the ConnectX® adapters. In case tuning is required, please refer to the <u>Performance Tuning for NVIDIA Adapters</u> Community post.

Features Overview and Configuration

i) 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
- <u>Storage Protocols</u>
- <u>Virtualization</u>

- <u>Resiliency</u>
- Docker Containers
- <u>HPC-X</u>
- Fast Driver Unload
- OVS Offload Using ASAP² Direct

Ethernet Network

The chapter contains the following sections:

- Ethernet Interface
- Quality of Service (QoS)
- <u>Ethtool</u>
- <u>Checksum Offload</u>
- Ignore Frame Check Sequence (FCS) Errors
- RDMA over Converged Ethernet (RoCE)
- Flow Control
- Explicit Congestion Notification (ECN)
- <u>RSS Support</u>
- <u>Time-Stamping</u>
- Flow Steering
- Wake-on-LAN (WoL)
- Hardware Accelerated 802.1ad VLAN (Q-in-Q Tunneling)
- VLAN Stripping in Linux Verbs
- Offloaded Traffic Sniffer

- <u>Dump Configuration</u>
- Local Loopback Disable
- Kernel Transport Layer Security (kTLS) Offloads
- IPsec Crypto Offload
- IPsec Full Offload
- MACsec Full Offload

Ethernet Interface

Port Type Management/VPI Cards Configuration

Ports of ConnectX-4 adapter cards and above can be individually configured to work as InfiniBand or Ethernet ports. By default, both VPI ports are initialized as InfiniBand ports. If you wish to change the port type, use the mlxconfig script after the driver is loaded.

For further information on how to set the port type, please refer to the MFT User Manual.

Counters

Counters are used to provide information about how well an operating system, an application, a service, or a driver is performing. The counter data help determine system bottlenecks and fine-tune the system and application performance. The operating system, network, and devices provide counter data that an application can consume to provide users with a graphical view of how well the system is performing.

The counter index is a Queue Pair (QP) attribute given in the QP context. Multiple QPs may be associated with the same counter set. If multiple QPs share the same counter, the counter value will represent the cumulative total.

RoCE Counters

- RoCE counters are available only through sysfs located under:
 - #/sys/class/infiniband/<device>/ports/*/hw_counters/
 - #/sys/class/infiniband/<device>/hw_counters/

• #/sys/class/infiniband/<device>/ports/*/counters/

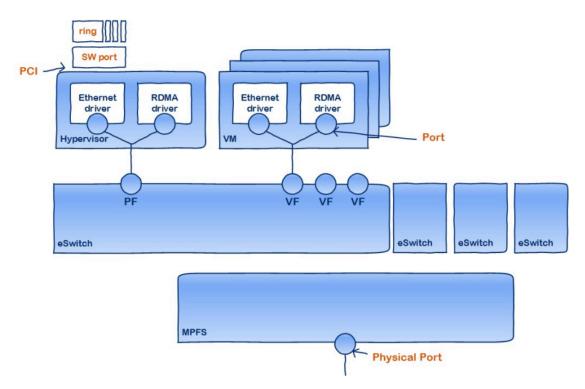
For mlx5 port and RoCE counters, refer to the <u>Understanding mlx5 Linux Counters</u> Community post.

SR-IOV Counters

Physical Function can also read Virtual Functions' port counters through sysfs located under # /sys/class/net/<interface_name>/device/sriov/<index>/stats/

ethtool Counters

The ethtool counters are counted in different places, according to which they are divided into groups. Each counters group may also have different counter types.



For the full list of supported ethtool counters, refer to the <u>Understanding mlx5 ethtool</u> <u>Counters</u> community post.

Persistent Naming

To avoid network interface renaming after boot or driver restart, create a file names "/etc/udev/rules.d/85-net-persistent-names.rules" with rules to set explicit interface names (the number, 85, must be more than 83).

• Example for Ethernet interfaces:

```
#PCI device 15b3:1019 (mlx5_core)
#NAME:="some name" , := is used to make sure that device name
will be persistent.
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*",
ATTR{address}=="00:02:c9:fa:c3:50", ATTR{dev_id}=="0x0",
ATTR{type}=="1", KERNEL=="eth*", NAME:="eth1"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*",
ATTR{address}=="00:02:c9:fa:c3:51", ATTR{dev_id}=="0x0",
ATTR{type}=="1", KERNEL=="eth*", NAME:="eth2"
```

• Example for IPoIB interfaces:

```
1. ATTR{type}=="32" is IPoIB interfaces
#2. Find address value for each IB interfaces:
#cat /sys/class/net/ib0/address | head -1 | cut -d":" -f"13-"
#24:8a:07:03:00:a5:29:38
#Use above output to place into ATTR{address}.
#Address match must start with ?* and only reference the last
8 bytes of the address
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{type}=="32",
ATTR{address}=="?*24:8a:07:03:00:a5:29:38", NAME:="ib0"
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{type}=="32",
ATTR{address}=="?*24:8a:07:03:00:a5:29:39", NAME:="ib1"
```

Interrupt Request (IRQ) Naming

Once IRQs are allocated by the driver, they are named mlx5_comp<x>@pci:<pci_addr>. The IRQ name is constant and is not affected by the interface state.

The mlx5_core driver allocates all IRQs during loading time to support the maximum possible number of channels. Once the driver is up, no further IRQs are freed or allocated. Changing the number of working channels does not re-allocate or free the IRQs.

Quality of Service (QoS)

Quality of Service (QoS) is a mechanism of assigning a priority to a network flow (socket, rdma_cm connection) and manage its guarantees, limitations and its priority over other flows. This is accomplished by mapping the user's priority to a hardware TC (traffic class) through a 2/3 stage process. The TC is assigned with the QoS attributes and the different flows behave accordingly.

Mapping Traffic to Traffic Classes

Mapping traffic to TCs consists of several actions which are user controllable, some controlled by the application itself and others by the system/network administrators.

The following is the general mapping traffic to Traffic Classes flow:

- 1. The application sets the required Type of Service (ToS).
- 2. The ToS is translated into a Socket Priority (sk_prio).
- 3. The sk_prio is mapped to a User Priority (UP) by the system administrator (some applications set sk_prio directly).
- 4. The UP is mapped to TC by the network/system administrator.
- 5. TCs hold the actual QoS parameters

QoS can be applied on the following types of traffic. However, the general QoS flow may vary among them:

- **Plain Ethernet** Applications use regular inet sockets and the traffic passes via the kernel Ethernet driver
- RoCE Applications use the RDMA API to transmit using Queue Pairs (QPs)
- Raw Ethernet QP Application use VERBs API to transmit using a Raw Ethernet QP

Plain Ethernet Quality of Service Mapping

Applications use regular inet sockets and the traffic passes via the kernel Ethernet driver. The following is the Plain Ethernet QoS mapping flow:

- 1. The application sets the ToS of the socket using setsockopt (IP_TOS, value).
- 2. ToS is translated into the sk_prio using a fixed translation:

```
TOS 0 <=> sk_prio 0
TOS 8 <=> sk_prio 2
TOS 24 <=> sk_prio 4
TOS 16 <=> sk_prio 6
```

- 3. The Socket Priority is mapped to the UP in the following conditions:
 - 1. If the underlying device is a VLAN device, egress_map is used controlled by the vconfig command. This is per VLAN mapping.
 - 2. If the underlying device is not a VLAN device, the mapping is done in the driver.
- 4. The UP is mapped to the TC as configured by the mlnx_qos tool or by the lldpad daemon if DCBX is used.

j) Note

Socket applications can use setsockopt (SK_PRIO, value) to directly set the sk_prio of the socket. In this case, the ToS to sk_prio fixed mapping is not needed. This allows the application and the administrator to utilize more than the 4 values possible via ToS.

i) Note

In the case of a VLAN interface, the UP obtained according to the above mapping is also used in the VLAN tag of the traffic.

RoCE Quality of Service Mapping

Applications use RDMA-CM API to create and use QPs. The following is the RoCE QoS mapping flow:

- 1. The application sets the ToS of the QP using the rdma_set_option option(RDMA_OPTION_ID_TOS, value).
- 2. ToS is translated into the Socket Priority (sk_prio) using a fixed translation:

```
TOS 0 <=> sk_prio 0
TOS 8 <=> sk_prio 2
TOS 24 <=> sk_prio 4
TOS 16 <=> sk_prio 6
```

- 3. The Socket Priority is mapped to the User Priority (UP) using the tc command.
- In the case of a VLAN device where the parent real device is used for the purpose of this mapping
- If the underlying device is a VLAN device, and the parent real device was not used for the mapping, the VLAN device's egress_map is used

4. UP is mapped to the TC as configured by the mlnx_qos tool or by the lldpad daemon if DCBX is used.

i) Note

With RoCE, there can only be 4 predefined ToS values for the purpose of QoS mapping.

Map Priorities with set_egress_map

For RoCE old kernels that do not support set_egress_map, use the tc_wrap script to map between sk_prio and UP. Use tc_wrap with option -u. For example:

```
tc_wrap -i <ethX> -u <skprio2up mapping>
```

Quality of Service Properties

The different QoS properties that can be assigned to a TC are:

- <u>Strict Priority</u>
- Enhanced Transmission Selection (ETS)
- <u>Rate Limit</u>
- <u>Trust State</u>
- <u>Receive Buffer</u>
- DCBX Control Mode

Strict Priority

When setting a TC's transmission algorithm to be 'strict', then this TC has absolute (strict) priority over other TC strict priorities coming before it (as determined by the TC number: TC 7 is the highest priority, TC 0 is lowest). It also has an absolute priority over nonstrict TCs (ETS).

This property needs to be used with care, as it may easily cause starvation of other TCs.

A higher strict priority TC is always given the first chance to transmit. Only if the highest strict priority TC has nothing more to transmit, will the next highest TC be considered.

Nonstrict priority TCs will be considered last to transmit.

This property is extremely useful for low latency low bandwidth traffic that needs to get immediate service when it exists, but is not of high volume to starve other transmitters in the system.

Enhanced Transmission Selection (ETS)

Enhanced Transmission Selection standard (ETS) exploits the time periods in which the offered load of a particular Traffic Class (TC) is less than its minimum allocated bandwidth by allowing the difference to be available to other traffic classes.

After servicing the strict priority TCs, the amount of bandwidth (BW) left on the wire may be split among other TCs according to a minimal guarantee policy.

If, for instance, TCO is set to 80% guarantee and TC1 to 20% (the TCs sum must be 100), then the BW left after servicing all strict priority TCs will be split according to this ratio.

Since this is a minimum guarantee, there is no maximum enforcement. This means, in the same example, that if TC1 did not use its share of 20%, the reminder will be used by TCO.

ETS is configured using the mlnx_qos tool (mlnx_qos) which allows you to:

- Assign a transmission algorithm to each TC (strict or ETS)
- Set minimal BW guarantee to ETS TCs

Usage:

mlnx_qos -i \[options\]

Rate Limit

Rate limit defines a maximum bandwidth allowed for a TC. Please note that 10% deviation from the requested values is considered acceptable.

Trust State

Trust state enables prioritizing sent/received packets based on packet fields.

The default trust state is PCP. Ethernet packets are prioritized based on the value of the field (PCP/DSCP).

For further information on how to configure Trust mode, please refer to <u>HowTo Configure</u> <u>Trust State on NVIDIA Adapters</u> community post.

i) Note

Setting the Trust State mode shall be done before enabling SR-IOV in order to propagate the Trust State to the VFs.

Receive Buffer

By default, the receive buffer configuration is controlled automatically. Users can override the receive buffer size and receive buffer's xon and xoff thresholds using mlnx_qos tool.

For further information, please refer to <u>HowTo Tune the Receive buffers on NVIDIA</u> <u>Adapters</u> community post.

DCBX Control Mode

DCBX settings, such as "ETS" and "strict priority" can be controlled by firmware or software. When DCBX is controlled by firmware, changes of QoS settings cannot be done by the software. The DCBX control mode is configured using the mlnx_qos -d os/fw command.

For further information on how to configure the DCBX control mode, please refer to <u>mlnx_qos</u> community post.

Quality of Service Tools

mlnx_qos

mlnx_qos is a centralized tool used to configure QoS features of the local host. It communicates directly with the driver thus does not require setting up a DCBX daemon on the system.

The mlnx_qos tool enables the administrator of the system to:

• Inspect the current QoS mappings and configuration

The tool will also display maps configured by TC and vconfig set_egress_map tools, in order to give a centralized view of all QoS mappings.

- Set UP to TC mapping
- Assign a transmission algorithm to each TC (strict or ETS)
- Set minimal BW guarantee to ETS TCs
- Set rate limit to TCs
- Set DCBX control mode
- Set cable length
- Set trust state

(i) Note

For an unlimited ratelimit, set the ratelimit to 0.

Usage

```
mlnx_qos -i <interface> \[options\]
```

Options

	Show the program's version number and exit
version	

-h, help	Show this help message and exit				
-f LIST, pfc=LIS T	Set priority flow control for each priority. LIST is a comma separated value for each priority starting from 0 to 7. Example: 0,0,0,0,1,1,1,1 enable PFC on TC4-7				
-p LIST, prio_tc =LIST	Maps UPs to TCs. LIST is 8 comma-separated TC numbers. Example: 0,0,0,0,1,1,1,1 maps UPs 0-3 to TC0, and UPs 4-7 to TC1				
-s LIST, tsa=LIS T	Transmission algorithm for each TC. LIST is comma separated algorithm names for each TC. Possible algorithms: strict, ets and vendor. Example: vendor,strict,ets,ets,ets,ets,ets,ets sets TCO to vendor, TC1 to strict, TC2-7 to ets				
-t LIST, tcbw=L IST	Set the minimally guaranteed %BW for ETS TCs. LIST is comma-separated percents for each TC. Values set to TCs that are not configured to ETS algorithm are ignored but must be present. Example: if TC0,TC2 are set to ETS, then 10,0,90,0,0,0,0,0,0,0 ill set TC0 to 10% and TC2 to 90%. Percents must sum to 100				
-r LIST, ratelimi t=LIST	Rate limit for TCs (in Gbps). LIST is a comma-separated Gbps limit for each TC. Example: 1,8,8 will limit TCO to 1Gbps, and TC1,TC2 to 8 Gbps each				
-d DCBX, - - dcbx=D CBX	Set dcbx mode to firmware controlled(fw) or OS controlled(os). Note, when in OS mode, mlnx_qos should not be used in parallel with other dcbx tools, such as lldptool				
 trust=T RUST	set priority trust state to pcp or dscp				
 dscp2p rio=DS CP2PRI O	Set/del a (dscp,prio) mapping. Example 'set,30,2' maps dscp 30 to priority 2. 'del,30,2' resets the dscp 30 mapping back to the default setting priority 0				
 cable_l en=CA	Set cable_len for buffer's xoff and xon thresholds				

BLE_LE N	
-i INTF, interfa ce=INT F	Interface name
-a	Show all interface's TCs

Get Current Configuration

```
ofed_scripts/utils/mlnx_qos -i ens1f0
DCBX mode: OS controlled
Priority trust state: dscp
dscp2prio mapping:
        prio:0 dscp:07,06,05,04,03,02,01,00,
        prio:1 dscp:15,14,13,12,11,10,09,08,
        prio:2 dscp:23,22,21,20,19,18,17,16,
        prio:3 dscp:31,30,29,28,27,26,25,24,
        prio:4 dscp:39,38,37,36,35,34,33,32,
        prio:5 dscp:47,46,45,44,43,42,41,40,
        prio:6 dscp:55,54,53,52,51,50,49,48,
        prio:7 dscp:63,62,61,60,59,58,57,56,
Cable len: 7
PFC configuration:
        priority 0 1 2 3 4 5 6 7
        enabled 0 0 0 0 0 0 0 0 0
tc: 0 ratelimit: unlimited, tsa: vendor
         priority: 1
tc: 1 ratelimit: unlimited, tsa: vendor
         priority: 0
tc: 2 ratelimit: unlimited, tsa: vendor
         priority: 2
tc: 3 ratelimit: unlimited, tsa: vendor
         priority: 3
tc: 4 ratelimit: unlimited, tsa: vendor
```

```
priority: 4
tc: 5 ratelimit: unlimited, tsa: vendor
priority: 5
tc: 6 ratelimit: unlimited, tsa: vendor
priority: 6
tc: 7 ratelimit: unlimited, tsa: vendor
priority: 7
```

Set ratelimit. 3Gbps for tc0 4Gbps for tc1 and 2Gbps for tc2

```
# mlnx_qos -i <interface> -p 0,1,2 -r 3,4,2
tc: 0 ratelimit: 3 Gbps, tsa: strict
         up: 0
                  skprio: 0
                  skprio: 1
                  skprio: 2 (tos: 8)
                  skprio: 3
                  skprio: 4 (tos: 24)
                  skprio: 5
                  skprio: 6 (tos: 16)
                  skprio: 7
                  skprio: 8
                  skprio: 9
                  skprio: 10
                  skprio: 11
                  skprio: 12
                  skprio: 13
                  skprio: 14
                  skprio: 15
         up: 3
         up: 4
         up: 5
         up: 6
         up: 7
```

ConfigureQoS. Map UP0,7 to tc0,1,2,3 to tc1 and 4,5,6 to tc2. Set tc0,tc1 as ets and tc2 as strict. Divide ets 30% for tc0 and 70% for tc1

```
# mlnx_qos -i <interface> -s ets,ets,strict -p 0,1,1,1,2,2,2 -t 30,70
tc: 0 ratelimit: 3 Gbps, tsa: ets, bw: 30%
         up: 0
                  skprio: 0
                  skprio: 1
                  skprio: 2 (tos: 8)
                  skprio: 3
                  skprio: 4 (tos: 24)
                  skprio: 5
                  skprio: 6 (tos: 16)
                  skprio: 7
                  skprio: 8
                  skprio: 9
                  skprio: 10
                  skprio: 11
                  skprio: 12
                  skprio: 13
                  skprio: 14
                  skprio: 15
  up: 7
tc: 1 ratelimit: 4 Gbps, tsa: ets, bw: 70%
         up: 1
         up: 2
         up: 3
tc: 2 ratelimit: 2 Gbps, tsa: strict
         up: 4
```

up: 5 up: 6

tc and tc_wrap.py

The tc tool is used to create 8 Traffic Classes (TCs).

The tool will either use the sysfs (/sys/class/net//qos/tc_num) or the tc tool to create the TCs.

Usage

```
tc_wrap.py -i <interface> \[options\]
```

Options

version	show program's version number and exit
-h,help	show this help message and exit
-u SKPRIO_UP, skprio_up=SKPRIO_UP	maps sk_prio to priority for RoCE. LIST is <=16 comma separated priority. index of element is sk_prio
-i INTF,interface=INTF	Interface name

Example

Run:

tc_wrap.py -i enp139s0

Output:

Tarrfic classes are set to 8

UP 0

```
skprio: 0 (vlan 5)
UP 1
        skprio: 1 (vlan 5)
UP 2
        skprio: 2 (vlan 5 tos: 8)
UP
   3
        skprio: 3 (vlan 5)
UP
    4
        skprio: 4 (vlan 5 tos: 24)
UP
    5
        skprio: 5 (vlan 5)
UP
    6
        skprio: 6 (vlan 5 tos: 16)
UP 7
        skprio: 7 (vlan 5)
```

Additional Tools

tc tool compiled with the sch_mqprio module is required to support kernel v2.6.32 or higher. This is a part of iproute2 package v2.6.32-19 or higher. Otherwise, an alternative custom sysfs interface is available.

- mlnx_qos tool (package: ofed-scripts) requires python version 2.5 < = X
- tc_wrap.py (package: ofed-scripts) requires python version 2.5 < = X

Packet Pacing

ConnectX-4 and above devices allow packet pacing (traffic shaping) per flow. This capability is achieved by mapping a flow to a dedicated send queue and setting a rate limit on that Send queue.

Note the following:

• Up to 512 send queues are supported

- 16 different rates are supported
- The rates can vary from 1 Mbps to line rate in 1 Mbps resolution
- Multiple queues can be mapped to the same rate (each queue is paced independently)
- It is possible to configure rate limit per CPU and per flow in parallel

System Requirements

- Driver v3.3 or higher
- Linux kernel v4.1 or higher
- ConnectX-4 or ConnectX-4 Lx adapter cards with an official firmware version

Packet Pacing Configuration

i Note

This configuration is non-persistent and does not survive driver restart.

1. Firmware Activation:

First, make sure MFT service (mst) is started:

```
# mst start
```

Then run:

#echo "MLNX_RAW_TLV_FILE" > /tmp/mlxconfig_raw.txt

```
#echo "0x0000004 0x0000010c 0x00000000 0x00000001" >>
/tmp/mlxconfig_raw.txt
#yes | mlxconfig -d <mst_dev> -f /tmp/mlxconfig_raw.txt
set_raw > /dev/null
#reboot /mlxfwreset
```

```
#echo "MLNX_RAW_TLV_FILE" > /tmp/mlxconfig_raw.txt
#echo "0x0000004 0x0000010c 0x00000000 0x00000000" >>
/tmp/mlxconfig_raw.txt
#yes | mlxconfig -d <mst_dev >-f /tmp/mlxconfig_raw.txt
set_raw > /dev/null
#reboot /mlxfwreset
```

2. Driver Activation:

There are two operation modes for Packet Pacing:

1. Rate limit per CPU core:

When XPS is enabled, traffic from a CPU core will be sent using the corresponding send queue. By limiting the rate on that queue, the transmit rate on that CPU core will be limited. For example:

```
echo 300 > /sys/class/net/ens2f1/queues/tx-0/tx_maxrate
```

In this case, the rate on Core 0 (tx-0) is limited to 300Mbit/sec.

- 2. Rate limit per flow:
 - 1. The driver allows opening up to 512 additional send queues using the following command:

ethtool -L ens2f1 other 1200

In this case, 1200 additional queues are opened

2. Create flow mapping.

Users can map a certain destination IP and/or destination layer 4 Port to a specific send queue. The match precedence is as follows:

- IP + L4 Port
- IP only
- L4 Port only
- No match (the flow would be mapped to default queues)

To create flow mapping:

Configure the destination IP. Write the IP address in hexadecimal representation to the relevant sysfs entry. For example, to map IP address 192.168.1.1 (0xc0a80101) to send queue 310, run the following command:

echo 0xc0a80101 > /sys/class/net/ens2f1/queues/tx-310/flow_map/dst_ip

To map Destination L4 3333 port (either TCP or UDP) to the same queue, run:

```
echo 3333 > /sys/class/net/ens2f1/queues/tx-
310/flow_map/dst_port
```

From this point on, all traffic destined to the given IP address and L4 port will be sent using send queue 310. All other traffic will be sent using the original send queue.

iii. Limit the rate of this flow using the following command:

echo 100 > /sys/class/net/ens2f1/queues/tx-310/tx_maxrate

i) Note

Each queue supports only a single IP+Port combination.

Ethtool

Ethtool is a standard Linux utility for controlling network drivers and hardware, particularly for wired Ethernet devices. It can be used to:

- Get identification and diagnostic information
- Get extended device statistics
- Control speed, duplex, auto-negotiation and flow control for Ethernet devices
- Control checksum offload and other hardware offload features
- Control DMA ring sizes and interrupt moderation
- Flash device firmware using a .mfa2 image

Ethtool Supported Options

Options	Description
ethtoolset-priv- flags eth <x> <priv flag> <on off=""></on></priv </x>	Enables/disables driver feature matching the given private flag.

Options	Description		
ethtoolshow-priv- flags eth <x></x>	Shows driver private flags and their states (ON/OFF).		
ethtool -a eth <x></x>	Queries the pause frame settings.		
ethtool -A eth <x> [rx on off] [tx on off]</x>	Sets the pause frame settings.		
ethtool -c eth <x></x>	Queries interrupt coalescing settings.		
ethtool -C eth <x> [pkt-rate-low N] [pkt- rate-high N] [rx- usecs-low N] [rx- usecs-high N]</x>	Sets the values for packet rate limits and for moderation time high and low values.		
ethtool -C eth <x> [rx- usecs N] [rx-frames N]</x>	Sets the interrupt coalescing setting. rx-frames will be enforced immediately, rx-usecs will be enforced only when adaptive moderation is disabled. Note : usec settings correspond to the time to wait after the *last* packet is sent/received before triggering an interrupt.		
ethtool -C eth <x> adaptive-rx on off</x>	Enables/disables adaptive interrupt moderation. By default, the driver uses adaptive interrupt moderation for the receive path, which adjusts the moderation time to the traffic pattern.		
ethtool -C eth <x> adaptive-tx on off</x>	Note : Supported by mlx5e for ConnectX-4 and above adapter cards. Enables/disables adaptive interrupt moderation. By default, the driver uses adaptive interrupt moderation for the transmit path, which adjusts the moderation parameters (time/frames) to the traffic pattern.		
ethtool -g eth <x></x>	Queries the ring size values.		
ethtool -G eth <x> [rx <n>] [tx <n>]</n></n></x>	Modifies the ring size.		
ethtool -i eth <x></x>	Checks driver and device information. For example:		
	driver: mlx5_core		
	version: 5.1-0.4.0		
	firmware-version: 4.6.4046 (MT_QEMU000000)		
	expansion-rom-version:		

Options	Description		
	bus-info: 0000:07:00.0		
	supports-statistics: yes		
	supports-test: yes		
	supports-eeprom-access: no		
	supports-register-dump: no		
	supports-priv-flags: yes		
ethtool -k eth <x></x>	Queries the stateless offload status.		
ethtool -K eth <x> [rx on off] [tx on off] [sg on off] [tso on off] [lro on off] [gro on off] [gso on off] [rxvlan on off] [txvlan on off] [ntuple on/off] [rxhash on/off] [rx-all on/off] [rx-fcs on/off]</x>	Sets the stateless offload status. TCP Segmentation Offload (TSO), Generic Segmentation Offload (GSO): increase outbound throughput by reducing CPU overhead. It works by queuing up large buffers and letting the network interface card split them into separate packets. Large Receive Offload (LRO): increases inbound throughput of high-bandwidth network connections by reducing CPU overhead. It works by aggregating multiple incoming packets from a single stream into a larger buffer before they are passed higher up the networking stack, thus reducing the number of packets that have to be processed. LRO is available in kernel versions < 3.1 for untagged traffic. Hardware VLAN insertion Offload (txvlan): When enabled, the sent VLAN tag will be inserted into the packet by the hardware. Note : LRO will be done whenever possible. Otherwise GRO will be done. Generic Receive Offload (GRO) is available throughout all kernels. Hardware VLAN Striping Offload (rxvlan): When enabled received VLAN traffic will be stripped from the VLAN tag by the hardware. RX FCS (rx-fcs): Keeps FCS field in the received packets.Sets the stateless offload status. RX FCS validation (rx-all): Ignores FCS validation on the received packets.		
ethtool -l eth <x></x>	Shows the number of channels.		
ethtool -L eth <x> [rx <n>] [tx <n>]</n></n></x>	Sets the number of channels. Notes : • This also resets the RSS table to its default distribution,		

Options	Description				
	 For ConnectX[®]-4 cards, use ethtool -L eth<x> combined <n> to set both RX and TX channels.</n></x> 				
ethtool -m dump- module-eeprom eth <x> [raw on off] [hex on off] [offset N] [length N]</x>	Queries/decodes the cable module eeprom information.				
ethtool -p identify DEVNAME	Enables visual identification of the port by LED blinking [TIME IN-SECONDS].				
ethtool -p identify eth <x> <led </led duration>Allows users to identify interface's physical port by turn ports LED on for a number of seconds.NoteNote: The limit for the LED duration is 65535 seconds.</x>					
ethtool -S eth <x></x>	Obtains additional device statistics.				
advertise <n> autoneg on</n>	Changes the advertised link modes to requested link modes <n> To check the link modes' hex values, run <man ethtool=""> and to check the supported link modes, run ethtool eth<x> For advertising new link modes, make sure to configure the entire bitmap as follows:</x></man></n>				
	200GAUI-4 / 200GBASE-CR4/KR4	0x7c00000000000000000000000000000000000			
	100GAUI-2 / 100GBASE-CR2 / KR2	0x3E000000000000000			
	CAUI-4 / 100GBASE-CR4 / KR4	0xF00000000			
	50GAUI-1 / LAUI-1/ 50GBASE-CR / KR	0x1F00000000000000			
	50GAUI-2 / LAUI-2/ 50GBASE- CR2/KR2	0x10C0000000			
	XLAUI-4/XLPPI-4 // 40G				
		0x7800000			
	25GAUI-1/25GBASE-CR/KR	0x7800000 0x38000000			
	25GAUI-1/ 25GBASE-CR / KR XFI / XAUI-1 // 10G				
		0x38000000			
	XFI / XAUI-1 // 10G	0x380000000 0x7C0000181000			

Options	Description		
	Notes:		
	 Both previous and new link modes configurations are supported, however, they must be run separately. Any link mode configuration on Kernels below v5.1 and ConnectX-6 HCAs will result in the advertisement of the full capabilities. <autoneg on=""> only sends a hint to the driver that the user wants to modify advertised link modes and not speed.</autoneg> 		
ethtool -s eth <x> msglvl [N]</x>	Changes the current driver message level.		
ethtool -s eth <x> speed <speed> autoneg off</speed></x>	Changes the link speed to requested <speed>. To check the supported speeds, run ethtool eth<x>. Note: does not set autoneg OFF, it only hints the driver to set a specific speed.</x></speed>		
ethtool -t eth <x></x>	Performs a self-diagnostics test.		
ethtool -T eth <x></x>	Shows time stamping capabilities		
ethtool -x eth <x></x>	Retrieves the receive flow hash indirection table.		
ethtool -X eth <x> equal a b c</x>	Sets the receive flow hash indirection table. Note : The RSS table configuration is reset whenever the number of channels is modified (using ethtool -L command).		
ethtoolshow-fec eth <x></x>	Queries current Forward Error Correction (FEC) encoding in case FEC is supported. Note : An output of "baser" implies Firecode encoding.		
ethtoolset-fec eth <x> encoding auto off rs baser</x>	Configures Forward Error Correction (FEC). Note : 'baser' encoding applies to the Firecode encoding, and 'auto' regards the HCA's default.		
ethtool -f flash <devname> FILE [N]</devname>	Flash firmware image on the device using the specified .mfa2 file (FILE). By default, the command flashes all the regions on the device unless a region number (N) is specified.		

Checksum Offload

The following Receive IP/L4 Checksum Offload modes are supported.

- **CHECKSUM_UNNECESSARY**: When this mode is used, the driver indicates to the Linux Networking Stack that the hardware successfully validated the IP and L4 checksum so the Linux Networking Stack does not need to deal with IP/L4 Checksum validation.
- **CHECKSUM_COMPLETE**: When this mode is used, the driver still reports to the OS the calculated by hardware checksum value. This allows accelerating checksum validation in Linux Networking Stack, since it does not have to calculate the whole checksum including payload by itself.
- **CHECKSUM_NONE**: When this mode is used, the driver indicates to the Linux Networking Stack that it must calculate and validate the IP/L4 checksum.

Ignore Frame Check Sequence (FCS) Errors

Upon receiving packets, the packets go through a checksum validation process for the FCS field. If the validation fails, the received packets are dropped.

When FCS is enabled (disabled by default), the device does not validate the FCS field even if the field is invalid.

It is not recommended to enable FCS.

For further information on how to enable/disable FCS, please refer to <u>ethtool option rx-fcs</u> <u>on/off</u>.

RDMA over Converged Ethernet (RoCE)

Remote Direct Memory Access (RDMA) is the remote memory management capability that allows server-to-server data movement directly between application memory without any CPU involvement. RDMA over Converged Ethernet (RoCE) is a mechanism to provide this efficient data transfer with very low latencies on lossless Ethernet networks. With advances in data center convergence over reliable Ethernet, ConnectX® Ethernet adapter cards family with RoCE uses the proven and efficient RDMA transport to provide the platform for deploying RDMA technology in mainstream data center application at 10GigE and 40GigE link-speed. ConnectX® Ethernet adapter cards family with its hardware offload support takes advantage of this efficient RDMA transport (InfiniBand) services over Ethernet to deliver ultra-low latency for performance-critical and transaction-intensive applications such as financial, database, storage, and content delivery networks.

When working with RDMA applications over Ethernet link layer the following points should be noted:

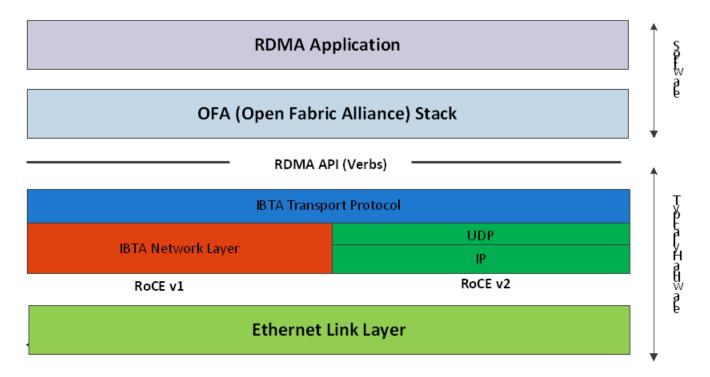
- The presence of a Subnet Manager (SM) is not required in the fabric. Thus, operations that require communication with the SM are managed in a different way in RoCE. This does not affect the API but only the actions such as joining the multicast group, that need to be taken when using the API
- Since LID is a layer 2 attribute of the InfiniBand protocol stack, it is not set for a port and is displayed as zero when querying the port
- With RoCE, the alternate path is not set for RC QP. Therefore, APM (another type of High Availability and part of the InfiniBand protocol) is not supported
- Since the SM is not present, querying a path is impossible. Therefore, the path record structure must be filled with relevant values before establishing a connection. Hence, it is recommended working with RDMA-CM to establish a connection as it takes care of filling the path record structure
- VLAN tagged Ethernet frames carry a 3-bit priority field. The value of this field is derived from the IB SL field by taking the 3 least significant bits of the SL field
- RoCE traffic is not shown in the associated Ethernet device's counters since it is offloaded by the hardware and does not go through Ethernet network driver. RoCE traffic is counted in the same place where InfiniBand traffic is counted; /sys/class/infiniband/<device>/ports/<port number>/counters/

RoCE Modes

RoCE encapsulates IB transport in one of the following Ethernet packets:

- RoCEv1 dedicated ether type (0x8915)
- **RoCEv2** UDP and dedicated UDP port (4791)

RoCEv1 and RoCEv2 Protocol Stack



RoCEv1

RoCE v1 protocol is defined as RDMA over Ethernet header (as shown in the figure above). It uses ethertype 0x8915 and can be used with or without the VLAN tag. The regular Ethernet MTU applies on the RoCE frame.

RoCEv2

A straightforward extension of the RoCE protocol enables traffic to operate in IP layer 3 environments. This capability is obtained via a simple modification of the RoCE packet format. Instead of the GRH used in RoCE, IP routable RoCE packets carry an IP header which allows traversal of IP L3 Routers and a UDP header (RoCEv2 only) that serves as a stateless encapsulation layer for the RDMA Transport Protocol Packets over IP.

The proposed RoCEv2 packets use a well-known UDP destination port value that unequivocally distinguishes the datagram. Similar to other protocols that use UDP encapsulation, the UDP source port field is used to carry an opaque flow-identifier that allows network devices to implement packet forwarding optimizations (e.g. ECMP) while staying agnostic to the specifics of the protocol header format.

Furthermore, since this change exclusively affects the packet format on the wire, and due to the fact that with RDMA semantics packets are generated and consumed below the AP, applications can seamlessly operate over any form of RDMA service, in a completely transparent way.

(i) Note

Both RoCEv1 and RoCEv2 are supported by default; the driver associates all GID indexes to RoCEv1 and RoCEv2, thus, a single entry for each RoCE version.

For further information, please refer to <u>Recommended Network</u> <u>Configuration Examples For RoCE Deployment</u> Community post.

GID Table Population

GID table entries are created whenever an IP address is configured on one of the Ethernet devices of the NIC's ports. Each entry in the GID table for RoCE ports has the following fields:

- GID value
- GID type
- Network device

The GID table is occupied with two GIDs, both with the same GID value but with different types. The network device in an entry is the Ethernet device with the IP address that GID is associated with. The GID format can be of 2 types; IPv4 and IPv6. IPv4 GID is an IPv4-mapped IPv6 address, while IPv6 GID is the IPv6 address itself. Layer 3 header for packets associated with IPv4 GIDs will be IPv4 (for RoCEv2) and IPv6/GRH for packets associated with IPv6 GIDs for RoCEv1.

GID table is exposed to userspace via sysfs

• GID values can be read from:

/sys/class/infiniband/{device}/ports/{port}/gids/{index}

• GID type can be read from:

/sys/class/infiniband/{device}/ports/{port}/gid_attrs/types/{indexisted

• GID net_device can be read from:

/sys/class/infiniband/{device}/ports/{port}/gid_attrs/ndevs/{indevice}/

Setting the RoCE Mode for a Queue Pair (QP)

Setting RoCE mode for devices that support two RoCE modes is different for RC/UC QPs (connected QP types) and UD QP.

To modify an RC/UC QP (connected QP) from INIT to RTR, an Address Vector (AV) must be given. The AV, among other attributes, should specify the index of the port's GID table for the source GID of the QP. The GID type in that index will be used to set the RoCE type of the QP.

Setting RoCE Mode of RDMA_CM Applications

RDMA_CM interface requires only the active side of the peer to pass the IP address of the passive side. The RDMA_CM decides upon the source GID to be used and obtains it from the GID table. Since more than one instance of the GID value is possible, the lookup should be also according to the GID type. The type to use for the lookup is defined as a global value of the RDMA_CM module. Changing the value of the GID type for the GID table lookups is done using the cma_roce_mode script.

• To print the current RoCE mode for a device port:

cma_roce_mode -d <dev> -p <port>

• To set the RoCE mode for a device port:

cma_roce_mode -d <dev> -p <port> -m <1|2>

GID Table Example

The following is an example of the GID table.

DEV	PORT	INDE X	GID	IPv4	Туре	Netdev
mlx5_ 0	1	0	fe80:0000:0000:0000:ba59:9fff:f e1a:e3ea		v1	p4p1
mlx5_ 0	1	1	fe80:0000:0000:0000:ba59:9fff:f e1a:e3ea		v2	p4p1
mlx5_ 0	1	2	0000:0000:0000:0000:0000:ffff:0 a0a:0a01	10.10.10.1	v1	p4p1
mlx5_ 0	1	3	0000:0000:0000:0000:0000:ffff:0 a0a:0a01	10.10.10.1	v2	p4p1
mlx5_ 1	1	0	fe80:0000:0000:0000:ba59:9fff:f e1a:e3eb		v1	p4p2
mlx5_ 1	1	1	fe80:0000:0000:0000:ba59:9fff:f e1a:e3eb		v2	p4p2

where:

- Entries on port 1 index 0/1 are the default GIDs, one for each supported RoCE type
- Entries on port 1 index 2/3 belong to IP address 192.168.1.70 on eth1
- Entries on port 1 index 4/5 belong to IP address 193.168.1.70 on eth1.100
- Packets from a QP that is associated with these GID indexes will have a VLAN header (VID=100)
- Entries on port 1 index 6/7 are IPv6 GID. Packets from a QP that is associated with these GID indexes will have an IPv6 header

RoCE Lossless Ethernet Configuration

In order to function reliably, RoCE requires a form of flow control. While it is possible to use global flow control, this is normally undesirable, for performance reasons.

The normal and optimal way to use RoCE is to use Priority Flow Control (PFC). To use PFC, it must be enabled on all endpoints and switches in the flow path.

Configuring SwitchX® Based Switch System

To enable RoCE, the SwitchX should be configured as follows:

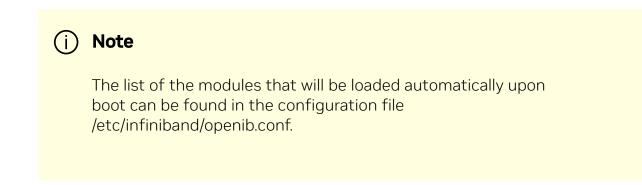
- Ports facing the host should be configured as access ports, and either use global pause or Port Control Protocol (PCP) for priority flow control
- Ports facing the network should be configured as trunk ports, and use Port Control Protocol (PCP) for priority flow control
- For further information on how to configure SwitchX, please refer to SwitchX User Manual

Installing and Loading the Driver

To install and load the driver:

1. Install MLNX_OFED (See Installation section for further details).

RoCE is installed as part of mlx5 and other modules upon driver's installation.



2. Query for the device's information. Example:

```
ibv_devinfo MLNX_OFED_LINUX-5.0-2.1.8.0:
```

3. Display the existing MLNX_OFED version.

```
ofed_info -s
hca_id: mlx5_0
                                           InfiniBand (0)
        transport:
        fw_ver:
                                           16.28.0578
        node_guid:
                                           ec0d:9a03:0044:3764
        sys_image_guid:
                                           ec0d:9a03:0044:3764
        vendor_id:
                                           0x02c9
        vendor_part_id:
                                           4121
        hw_ver:
                                           0x0
        board_id:
                                           MT_000000009
        phys_port_cnt:
                                           1
                 port:
                          1
                                                    PORT_ACTIVE
                          state:
(4)
                                                    4096 (5)
                          max_mtu:
                          active_mtu:
                                                    1024(3)
                          sm_lid:
                                                    0
                          port_lid:
                                                    0
                          port_lmc:
                                                    0x00
                          link_layer:
                                                    Ethernet
```

Output Notes:

The port's state is: Ethernet is in PORT_ACTIVE state	The port state can also be obtained by running the following command: # cat /sys/class/infiniband/mlx5_0/ports/1/state: ACTIVE		
link_layer parameter shows that port 1 is Ethernet	The link_layer can also be obtained by running the following command: # cat /sys/class/infiniband/mlx5_0/ports/1/link_layer Ethernet		

The fw_ver parameter shows that the firmware version is 16.28.0578.

The firmware version can also be obtained by running the following command: # cat /sys/class/infiniband/mlx5_0/fw_ver 16.28.0578

Associating InfiniBand Ports to Ethernet Ports

The mlx5_ib driver holds a reference to the net device for getting notifications about the state of the port, as well as using the mlx5_core driver to resolve IP addresses to MAC that are required for address vector creation. However, RoCE traffic does not go through the mlx5_core driver; it is completely offloaded by the hardware.

```
# ibdev2netdev
mlx5_0 port 1 <===> eth2
#
```

Configuring an IP Address to the netdev Interface

To configure an IP address to the netdev interface:

1. Configure an IP address to the netdev interface on both sides of the link.

```
# ifconfig eth2 20.4.3.220
# ifconfig eth2
eth2 Link encap:Ethernet HWaddr 00:02:C9:08:E8:11
    inet addr:20.4.3.220 Bcast:20.255.255.Mask:255.0.0.0
    UP BROADCAST MULTICAST MTU:1500 Metric:1
    RX packets:0 errors:0 dropped:0 overruns:0 frame:0
    TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

2. Make sure that ping is working.

```
ping 20.4.3.219
PING 20.4.3.219 (20.4.3.219) 56(84) bytes of data.
64 bytes from 20.4.3.219: icmp_seq=1 ttl=64 time=0.873 ms
64 bytes from 20.4.3.219: icmp_seq=2 ttl=64 time=0.198 ms
64 bytes from 20.4.3.219: icmp_seq=3 ttl=64 time=0.167 ms
20.4.3.219 ping statistics -
3 packets transmitted, 3 received, 0% packet loss, time
2000ms rtt min/avg/max/mdev = 0.167/0.412/0.873/0.326 ms
```

Adding VLANs

To add VLANs:

1. Make sure that the 8021.q module is loaded.

modprobe 8021q

2. Add VLAN.

```
# vconfig add eth2 7
Added VLAN with VID == 7 to IF -:eth2:-
#
```

3. Configure an IP address.

ifconfig eth2.7 7.4.3.220

Defining Ethernet Priority (PCP in 802.1q Headers)

1. Define Ethernet priority on the server.

```
# ibv_rc_pingpong -g 1 -i 2 -l 4
local address: LID 0x0000, QPN 0x1c004f, PSN 0x9daf6c, GID
fe80::202:c900:708:e799
remote address: LID 0x0000, QPN 0x1c004f, PSN 0xb0a49b, GID
fe80::202:c900:708:e811
8192000 bytes in 0.01 seconds = 4840.89 Mbit/sec
1000 iters in 0.01 seconds = 13.54 usec/iter
```

2. Define Ethernet priority on the client.

```
# ibv_rc_pingpong -g 1 -i 2 -l 4 sw419
local address: LID 0x0000, QPN 0x1c004f, PSN 0xb0a49b, GID
fe80::202:c900:708:e811
remote address: LID 0x0000, QPN 0x1c004f, PSN 0x9daf6c, GID
fe80::202:c900:708:e799
8192000 bytes in 0.01 seconds = 4855.96 Mbit/sec
1000 iters in 0.01 seconds = 13.50 usec/iter
```

Using rdma_cm Tests

1. Use rdma_cm test on the server.

```
# ucmatose
cmatose: starting server
initiating data transfers
completing sends
```

receiving data transfers
data transfers complete
cmatose: disconnecting
disconnected
test complete
return status 0
#

2. Use rdma_cm test on the client.

```
# ucmatose -s 20.4.3.219
cmatose: starting client
cmatose: connecting
receiving data transfers
sending replies
data transfers complete
test complete
return status 0
#
```

This server-client run is without PCP or VLAN because the IP address used does not belong to a VLAN interface. If you specify a VLAN IP address, then the traffic should go over VLAN.

Type Of Service (ToS)

The TOS field for rdma_cm sockets can be set using the rdma_set_option() API, just as it is set for regular sockets. If a TOS is not set, the default value (0) is used. Within the rdma_cm kernel driver, the TOS field is converted into an SL field. The conversion formula is as follows:

• SL = TOS >> 5 (e.g., take the 3 most significant bits of the TOS field)

In the hardware driver, the SL field is converted into PCP by the following formula:

• PCP = SL & 7 (take the 3 least significant bits of the TOS field)

i) Note

SL affects the PCP only when the traffic goes over tagged VLAN frames.

DSCP

A new entry has been added to the RDMA-CM configfs that allows users to select default TOS for RDMA-CM QPs. This is useful for users that want to control the TOS field without changing their code. Other applications that set the TOS explicitly using the rdma_set_option API will continue to work as expected to override the configfs value.

For further information about DSCP marking, refer to <u>HowTo Set Egress ToS/DSCP on</u> <u>RDMA CM QPs</u> Community post.

RoCE LAG

RoCE LAG is a feature meant for mimicking Ethernet bonding for IB devices and is available for dual port cards only.

This feature is supported on kernel versions 4.9 and above.

RoCE LAG mode is entered when both Ethernet interfaces are configured as a bond in one of the following modes:

- active-backup (mode 1)
- balance-xor (mode 2)
- 802.3ad (LACP) (mode 4)

Any change of bonding configuration that negates one of the above rules (i.e, bonding mode is not 1, 2 or 4, or both Ethernet interfaces that belong to the same card are not the only slaves

of the bond interface), will result in exiting RoCE LAG mode and the return to normal IB device per port configuration.

Once RoCE LAG is enabled, instead of having two IB devices; mlx5_0 and mlx5_1, there will be one device named mlx5_bond_0.

For information on how to configure RoCE LAG, refer to <u>HowTo Configure RoCE over LAG</u> (<u>ConnectX-4/ConnectX-5/ConnectX-6</u>) Community post.

Disabling RoCE

By default, RoCE is enabled on all mlx5 devices. When RoCE is enabled, all traffic to UDP port 4791 is treated as RoCE traffic by the device.

In case you are only interested in Ethernet (no RDMA) and wish to enable forwarding of traffic to this port, you can disable RoCE through sysfs:

```
echo <0|1> > /sys/devices/{pci-bus-address}/roce_enable
```

j) Note

Once RoCE is disabled, only Ethernet traffic will be supported. Therefore, there will be no GID tables and only Raw Ethernet QPs will be supported.

The current RoCE state can be queried by sysfs:

cat /sys/devices/{pci-bus-address}/roce_enable

Enabling/Disabling RoCE on VMs via VFs

By default, when configuring VFs on the hypervisor, all VFs will be enabled with RoCE. This means they require more OS memory (from the VM). In case you are only interested in

Ethernet (no RDMA) on the VM, and you wish to save the VM memory, you can disable RoCE on the VF from the hypervisor. In addition, by disabling RoCE, a VM can have the capability of utilizing the RoCE UDP port (4791) for standard UDP traffic.

For details on how to enable/disable RoCE on a VF, refer to <u>HowTo Enable/Disable RoCE on</u> <u>VMs via VFs</u> Community post.

Force DSCP

This feature enables setting a global traffic_class value for all RC QPs, or setting a specific traffic class based on several matching criteria.

Usage

• To set a single global traffic class to be applied to all QPs, write the desired global traffic_class value to /sys/class/infiniband/<dev>/tc/<port>/traffic_class.

Note the following:

- Negative values indicate that the feature is disabled. traffic_class value can be set using ibv_modify_qp()
- Valid values range between 0 255

i) Note

The ToS field is 8 bits, while the DSCP field is 6 bits. To set a DSCP value of X, you need to multiply this value by 4 (SHIFT 2). For example, to set DSCP value of 24, set the ToS bit to 96 (24x4=96).

• To set multiple traffic class values based on source and/or destination IPs, write the desired rule to /sys/class/infiniband/<dev>/tc/<port>/traffic_class. For example:

```
echo "tclass=16,src_ip=1.1.1.2,dst_ip=1.1.1.0/24" >
/sys/class/infiniband/mlx5_0/tc/1/traffic_class
```

Note: Adding "tclass" prefix to tclass value is optional.

In the example above, traffic class 16 will be set to any QP with source IP 1.1.1.2 and destination IP 1.1.1.0/24.

Note that when setting a specific traffic class, the following rule precedence will apply:

- If a global traffic class value is set, it will be applied to all QPs
- If no global traffic class value is set, and there is a rule with matching source and destination IPs applicable to at least one QP, it will be applied
- Rules only with matching source and/or destination IPs have no defined precedence over other rules with matching source and/or destination IPs

Notes:

- A mask can be provided when using destination IPv4 addresses
- The rule precedence is not affected by the order in which rules are inserted
- Overlapping rules are entirely up to the administrator.
- "tclass=-1" will remove the rule from the database

Force Time to Live (TTL)

This feature enables setting a global TTL value for all RC QPs.

Write the desired TTL value to /sys/class/infiniband/<dev>/tc/<port>/ttl. Valid values range between 0 - 255

Flow Control

Priority Flow Control (PFC)

Priority Flow Control (PFC) IEEE 802.1Qbb applies pause functionality to specific classes of traffic on the Ethernet link. For example, PFC can provide lossless service for the RoCE traffic and best-effort service for the standard Ethernet traffic. PFC can provide different levels of service to specific classes of Ethernet traffic (using IEEE 802.1p traffic classes).

Configuring PFC on ConnectX-4 and above

1. Enable PFC on the desired priority:

```
mlnx_qos -i <ethX> --pfc <0/1>,<0/1>,<0/1>,<0/1>,<0/1>,<0/1>,<0/1>,
```

Example (Priority=4):

mlnx_qos -i eth1 --pfc 0,0,0,0,1,0,0,0

2. Create a VLAN interface:

vconfig add <ethX> <VLAN_ID>

Example (VLAN_ID=5):

vconfig add eth1 5

- 3. Set egress mapping:
 - 1. For Ethernet traffic:

vconfig set_egress_map <vlan_einterface> <skprio> <up>

Example (skprio=3, up=5):

vconfig set_egress_map eth1.5 3 5

4. Create 8 Traffic Classes (TCs):

```
tc_wrap.py -i <interface>
```

5. Enable PFC on the switch.

For information on how to enable PFC on your respective switch, please refer to Switch FC/PFC Configuration sections in the <u>RDMA/RoCE Solutions</u> Community page .

PFC Configuration Using LLDP DCBX

PFC Configuration on Hosts

PFC Auto-Configuration Using LLDP Tool in the OS

1. Start lldpad daemon on host.

lldpad -d

OR

service lldpad start

2. Send lldpad packets to the switch.

```
lldptool set-lldp -i <ethX> adminStatus=rxtx
lldptool -T -i <ethX> -V sysName enableTx=yess
lldptool -T -i <ethX> -V portDesc enableTx=yess
lldptool -T -i <ethX> -V sysDesc enableTx=yess
lldptool -T -i <ethX> -V sysCap enableTx=yess
lldptool -T -i <ethX> -V mngAddr enableTx=yess
```

```
lldptool -T -i <ethX> -V PFC enableTx=yes;
lldptool -T -I <ethX> -V CEE-DCBX enableTx=yess
```

- 3. Set the PFC parameters.
- For the CEE protocol, use dcbtool:

dcbtool sc <ethX> pfc pfcup:<xxxxxxx>

Example:

dcbtool sc eth6 pfc pfcup:01110001

where:

[nfound)	Enables/disables priority flow control. From left to right (priorities 0-7) - x
	can be equal to either 0 or 1. 1 indicates that the priority is configured
******	to transmit priority pause.

• For IEEE protocol, use lldptool:

```
lldptool -T -i <ethX> -V PFC enabled=x,x,x,x,x,x,x,x,x
```

Example:

lldptool -T -i eth2 -V PFC enabled=1,2,4

where:

PFC Auto-Configuration Using LLDP in the Firmware (for mlx5 driver)

There are two ways to configure PFC and ETS on the server:

- 1. Local Configuration Configuring each server manually.
- 2. **Remote Configuration** Configuring PFC and ETS on the switch, after which the switch will pass the configuration to the server using LLDP DCBX TLVs.

There are two ways to implement the remote configuration using mlx5 driver:

- 1. Configuring the adapter firmware to enable DCBX.
- 2. Configuring the host to enable DCBX.

For further information on how to auto-configure PFC using LLDP in the firmware, refer to the <u>HowTo Auto-Config PFC and ETS on ConnectX-4 via LLDP DCBX</u> Community post.

PFC Configuration on Switches

1. In order to enable DCBX, LLDP should first be enabled:

```
switch (config) # lldp
show lldp interfaces ethernet remote
```

2. Add DCBX to the list of supported TLVs per required interface.

For IEEE DCBX:

switch (config) # interface 1/1

switch (config interface ethernet 1/1) # lldp tlv-select dcbx

For CEE DCBX:

```
switch (config) # interface 1/1
switch (config interface ethernet 1/1) # lldp tlv-select dcbx-cee
```

3. [**Optional**] Application Priority can be configured on the switch, with the required ethertype and priority. For example, IP packet, priority 1:

```
switch (config) # dcb application-priority 0x8100 1
```

- 4. Make sure PFC is enabled on the host (for enabling PFC on the host, refer to <u>PFC</u> <u>Configuration on Hosts</u> section above). Once it is enabled, it will be passed in the LLDP TLVs.
- 5. Enable PFC with the desired priority on the Ethernet port.

```
dcb priority-flow-control enable force
dcb priority-flow-control priority <priority> enable
interface ethernet <port> dcb priority-flow-control mode on
force
```

Example - Enabling PFC with priority 3 on port 1/1:

```
dcb priority-flow-control enable force
dcb priority-flow-control priority 3 enable
interface ethernet 1/1 dcb priority-flow-control mode on force
```

Priority Counters

Several ingress and egress counters per priority are supported. Run ethtool -S to get the full list of port counters.

ConnectX-4 Counters

- Rx and Tx Counters:
 - Packets
 - Bytes
 - Octets
 - Frames
 - Pause
 - Pause frames
 - Pause Duration
 - Pause Transition

Example

```
# ethtool -S eth35 | grep prio4
    prio4_rx_octets: 62147780800
    prio4_rx_frames: 14885696
    prio4_tx_octets: 0
    prio4_tx_frames: 0
    prio4_rx_pause: 0
    prio4_rx_pause_duration: 0
    prio4_tx_pause: 26832
    prio4_tx_pause_duration: 14508
    prio4_rx_pause_transition: 0
```

The Pause counters in ConnectX-4 are visible via ethtool only for	(j) Note
phonties on which PFC is enabled.		The Pause counters in ConnectX-4 are visible via ethtool only for priorities on which PFC is enabled.

PFC Storm Prevention

PFC storm prevention enables toggling between default and auto modes.

The stall prevention timeout is configured to 8 seconds by default. Auto mode sets the stall prevention timeout to be 100 msec.

The feature can be controlled using sysfs in the following directory: /sys/class/net/eth*/settings/ pfc_stall_prevention

• To query the PFC stall prevention mode:

cat /sys/class/net/eth*/settings/pfc_stall_prevention

Example

\$ cat /sys/class/net/ens6/settings/pfc_stall_prevention
default

• To configure the PFC stall prevention mode:

echo <option>
/sys/class/net/<interface>/settings/pfc_stall_prevention

The following two counters were added to the ethtool -S:

- tx_pause_storm_warning_events when the device is stalled for a period longer than a pre-configured watermark, the counter increases, allowing the debug utility an insight into current device status.
- tx_pause_storm_error_events when the device is stalled for a period longer than a pre-configured timeout, the pause transmission is disabled, and the counter increase.

Dropless Receive Queue (RQ)

Dropless RQ feature enables the driver to notify the FW when SW receive queues are overloaded. This scenario takes place when the handling of SW receive queue is slower than the handling of the HW receive queues.

When this feature is enabled, a packet that is received while the receive queue is full will not be immediately dropped. The FW will accumulate these packets assuming posting of new WQEs will resume shortly. If received WQEs are not posted after a certain period of time, out_of_buffer counter will increase, indicating that the packet has been dropped.

This feature is disabled by default. In order to activate it, ensure that Flow Control feature is also enabled.

```
ethtool --set-priv-flags ens6 dropless_rq on
```

To get the feature state, run:

ethtool --show-priv-flags DEVNAME

Output example:

```
Private flags for DEVNAME:
rx_cge_moder : on
```

```
rx_cqe_compress: off
sniffer : off
dropless_rq : off
hw_lro : off
```

To disable the feature, run:

ethtool --set-priv-flags ens6 dropless_rq off

Explicit Congestion Notification (ECN)

ECN is an extension to the IP protocol. It allows reliable communication by notifying all ends of communication when congestion occurs. This is done without dropping packets.

Please note that this feature requires all nodes in the path (nodes, routers etc) between the communicating nodes to support ECN to ensure reliable communication. ECN is marked as 2 bits in the traffic control IP header. This ECN implementation refers to RoCE v2.

Enabling ECN

To enable ECN on the hosts:

1. Enable ECN in sysfs.

/sys/class/net/<interface>/ecn/<protocol>/enable/<X>

2. Query the attribute.

cat /sys/class/net/<interface>/ecn/<protocol>/params/<requested
attribute>

3. Modify the attribute.

echo <value>
/sys/class/net/<interface>/ecn/<protocol>/params/<requested
attribute>

ECN supports the following algorithms:

- r_roce_ecn_rp Reaction point
- r_roce_ecn_np Notification point

Each algorithm has a set of relevant parameters and statistics, which are defined per device, per port, per priority.

cat /sys/class/net/<interface>/ecn/<protocol>/enable/X

cat /sys/class/net/<interface>/ecn/<protocol>/requested attributes

echo 1 > /sys/class/net/<interface>/ecn/<protocol>/enable/X

echo <value> > /sys/class/net/<interface>/ecn/<protocol>/requested
attributes

where:

- X: priority {0..7}
- protocol: roce_rp / roce_np

• requested attributes: Next Slide for each protocol.

RSS Support

The device has the ability to use XOR as the RSS distribution function, instead of the default Toplitz function.

The XOR function can be better distributed among driver's receive queues in a small number of streams, where it distributes each TCP/UDP stream to a different queue.MLNX_OFEDMLNX_EN provides the following option to change the working RSS hash function from Toplitz to XOR, and vice-versa:

Through sysfs, located at: /sys/class/net/eth*/settings/hfunc.

To query the operational and supported hash functions:

cat /sys/class/net/eth*/settings/hfunc

Example:

cat /sys/class/net/eth2/settings/hfunc Operational hfunc: toeplitz Supported hfuncs: xor toeplitz

To change the operational hash function:

≻

echo xor > /sys/class/net/eth*/settings/hfunc

RSS Hash Function

The device has the ability to use XOR as the RSS distribution function, instead of the default Toplitz function.

The XOR function can be better distributed among driver's receive queues in a small number of streams, where it distributes each TCP/UDP stream to a different queue.MLNX_OFEDMLNX_EN provides the following option to change the working RSS hash function from Toplitz to XOR, and vice-versa:

Through sysfs, located at: /sys/class/net/eth*/settings/hfunc.

> To query the operational and supported hash functions:

cat /sys/class/net/eth*/settings/hfunc

Example:

cat /sys/class/net/eth2/settings/hfunc Operational hfunc: toeplitz Supported hfuncs: xor toeplitz

To change the operational hash function:

≻

echo xor > /sys/class/net/eth*/settings/hfunc

RSS Verbs Support

Receive Side Scaling (RSS) technology allows spreading incoming traffic between different receive descriptor queues. Assigning each queue to different CPU cores allows better load balancing of the incoming traffic and improves performance.

This technology was extended to user space by the verbs layer and can be used for RAW ETH QP.

RSS Flow Steering

Steering rules classify incoming packets and deliver a specific traffic type (e.g. TCP/UDP, IP only) or a specific flow to "RX Hash" QP. "RX Hash" QP is responsible for spreading the traffic it handles between the Receive Work Queues using RX hash and Indirection Table. The Receive Work Queue can point to different CQs that can be associated with different CPU cores.

Verbs

The below verbs should be used to achieve this task in both control and data path. Details per verb should be referenced from its man page.

- ibv_create_wq, ibv_modify_wq, ibv_destory_wq
- ibv_create_rwq_ind_table, ibv_destroy_rwq_ind_table
- ibv_create_qp_ex with specific RX configuration to create the "RX hash" QP

Time-Stamping

Time-Stamping Service

Time-stamping is the process of keeping track of the creation of a packet. A timestamping service supports assertions of proof that a datum existed before a particular time. Incoming packets are time-stamped before they are distributed on the PCI depending on the congestion in the PCI buffers. Outgoing packets are time-stamped very close to placing them on the wire.

Enabling Time-Stamping

Time-stamping is off by default and should be enabled before use.

To enable time-stamping for a socket:

Call setsockopt() with SO_TIMESTAMPING and with the following flags:

SOF_TIMESTAMPING_TX_HAR DWARE:	try to obtain send time-stamp in hardware
SOF_TIMESTAMPING_TX_SOF TWARE:	if SOF_TIMESTAMPING_TX_HARDWARE is off or fails, then do it in software
SOF_TIMESTAMPING_RX_HAR DWARE:	return the original, unmodified time-stamp as generated by the hardware
SOF_TIMESTAMPING_RX_SOF TWARE:	if SOF_TIMESTAMPING_RX_HARDWARE is off or fails, then do it in software
SOF_TIMESTAMPING_RAW_H ARDWARE:	return original raw hardware time-stamp
SOF_TIMESTAMPING_SYS_HA RDWARE:	return hardware time-stamp transformed into the system time base
SOF_TIMESTAMPING_SOFTW ARE:	return system time-stamp generated in software
SOF_TIMESTAMPING_TX/RX	determine how time-stamps are generated
SOF_TIMESTAMPING_RAW/SY S	determine how they are reported

*To enable time-stamping fo*r a net dev*ice:*

Admin privileged user can enable/disable time stamping through calling ioctl (sock, SIOCSH-WTSTAMP, &ifreq) with the following values:

• Send side time sampling, enabled by ifreq.hwtstamp_config.tx_type when:

```
/* possible values for hwtstamp_config->tx_type */
enum hwtstamp_tx_types {
    /*
    * No outgoing packet will need hardware time stamping;
    * should a packet arrive which asks for it, no hardware
    * time stamping will be done.
    /*
    HWTSTAMP_TX_OFF,
    /*
```

```
* Enables hardware time stamping for outgoing packets;
    * the sender of the packet decides which are to be
    * time stamped by setting %SOF_TIMESTAMPING_TX_SOFTWARE
    * before sending the packet.
    */
           HWTSTAMP_TX_ON,
/*
    * Enables time stamping for outgoing packets just as
    * HWTSTAMP_TX_ON does, but also enables time stamp insertion
    * directly into Sync packets. In this case, transmitted Sync
    * packets will not received a time stamp via the socket error
    * queue.
    */
           HWTSTAMP_TX_ONESTEP_SYNC,
};
Note: for send side time stamping currently only
HWTSTAMP_TX_OFF and
HWTSTAMP_TX_ON are supported.
```

• Receive side time sampling, enabled by ifreq.hwtstamp_config.rx_filter when:

```
/* possible values for hwtstamp_config->rx_filter */
enum hwtstamp_rx_filters {
    /* time stamp no incoming packet at all */
    HWTSTAMP_FILTER_NONE,
    /* time stamp any incoming packet */
    HWTSTAMP_FILTER_ALL,
    /* return value: time stamp all packets requested plus some others */
    HWTSTAMP_FILTER_SOME,
    /* PTP v1, UDP, any kind of event packet */
    HWTSTAMP_FILTER_PTP_V1_L4_EVENT,
    /* PTP v1, UDP, Sync packet */
    HWTSTAMP_FILTER_PTP_V1_L4_SYNC,
    /* PTP v1, UDP, Delay_req packet */
    HWTSTAMP_FILTER_PTP_V1_L4_DELAY_REQ,
```

```
/* PTP v2, UDP, any kind of event packet */
          HWTSTAMP_FILTER_PTP_V2_L4_EVENT,
          /* PTP v2, UDP, Sync packet */
          HWTSTAMP_FILTER_PTP_V2_L4_SYNC,
          /* PTP v2, UDP, Delay_req packet */
          HWTSTAMP_FILTER_PTP_V2_L4_DELAY_REQ.
          /* 802.AS1, Ethernet, any kind of event packet */
          HWTSTAMP_FILTER_PTP_V2_L2_EVENT,
          /* 802.AS1, Ethernet, Sync packet */
          HWTSTAMP_FILTER_PTP_V2_L2_SYNC,
          /* 802.AS1, Ethernet, Delay_reg packet */
          HWTSTAMP_FILTER_PTP_V2_L2_DELAY_REQ,
          /* PTP v2/802.AS1, any layer, any kind of event packet */
          HWTSTAMP_FILTER_PTP_V2_EVENT,
          /* PTP v2/802.AS1, any layer, Sync packet */
          HWTSTAMP_FILTER_PTP_V2_SYNC,
          /* PTP v2/802.AS1, any layer, Delay_reg packet */
          HWTSTAMP_FILTER_PTP_V2_DELAY_REQ,
};
Note: for receive side time stamping currently only
HWTSTAMP_FILTER_NONE and
HWTSTAMP_FILTER_ALL are supported.
```

Getting Time-Stamping

Once time stamping is enabled time stamp is placed in the socket Ancillary data. recvmsg() can be used to get this control message for regular incoming packets. For send time stamps the outgoing packet is looped back to the socket's error queue with the send time-stamp(s) attached. It can

be received with recvmsg (flags=MSG_ERRQUEUE). The call returns the original outgoing packet data including all headers prepended down to and including the link layer, the scm_time-stamping control message and a sock_extended_err control message with ee_errno==ENOMSG and ee_origin==SO_EE_ORIGIN_TIMESTAMPING. A socket with such

a pending bounced packet is ready for reading as far as select() is concerned. If the outgoing packet has to be fragmented, then only the first fragment is time stamped and returned to the sending socket.

(i) Note

When time-stamping is enabled, VLAN stripping is disabled. For more info please refer to Documentation/networking/timestamping.txt in <u>kernel.org</u>.



On ConnectX-4 and above adapter cards, when time-stamping is enabled, RX CQE compression is disabled (features are mutually exclusive).

Time Stamping Capabilities via ethtool

To display Time Stamping capabilities via ethtool:

Show Time Stamping capabilities:

ethtool -T eth<x>

Example:

ethtool -T eth0 Time stamping parameters for p2p1: Capabilities:

```
hardware-transmit
(SOF_TIMESTAMPING_TX_HARDWARE)
software-transmit
(SOF_TIMESTAMPING_TX_SOFTWARE)
hardware-receive
(SOF_TIMESTAMPING_RX_HARDWARE)
software-receive
(SOF_TIMESTAMPING_RX_SOFTWARE)
software-system-clock
                                 (SOF_TIMESTAMPING_SOFTWARE)
hardware-raw-clock
(SOF_TIMESTAMPING_RAW_HARDWARE)
PTP Hardware Clock: 1
Hardware Transmit Timestamp Modes:
off
(HWTSTAMP_TX_OFF)
on
(HWTSTAMP_TX_ON)
Hardware Receive Filter Modes:
none
(HWTSTAMP_FILTER_NONE)
all
(HWTSTAMP_FILTER_ALL)
```

For more details on PTP Hardware Clock, please refer to: <u>https://www.kernel.org/doc/Documentation/ptp/ptp.txt</u>

Steering PTP Traffic to Single RX Ring

As a result of Receive Side Steering (RSS) PTP traffic coming to UDP ports 319 and 320, it may reach the user space application in an out of order manner. In order to prevent this, PTP traffic needs to be steered to single RX ring using ethtool.

Example:

```
# ethtool -u ens7
8 RX rings available
Total 0 rules
# ethtool -U ens7 flow-type udp4 dst-port 319 action 0 loc 1
# ethtool -U ens7 flow-type udp4 dst-port 320 action 0 loc 0
# ethtool -u ens7
8 RX rings available
Total 2 rules
Filter: 0
Rule Type: UDP over IPv4
Src IP addr: 0.0.00 mask: 255,255,255,255
Dest IP addr: 0.0.0.0 mask: 255.255.255.255
TOS: 0x0 mask: 0xff
Src port: 0 mask: 0xffff
Dest port: 320 mask: 0x0
Action: Direct to queue 0
Filter: 1
Rule Type: UDP over IPv4
Src IP addr: 0.0.0.0 mask: 255.255.255.255
Dest IP addr: 0.0.00 mask: 255.255.255.255
TOS: 0x0 mask: 0xff
Src port: 0 mask: 0xffff
Dest port: 319 mask: 0x0
Action: Direct to queue 0
```

Tx Port Time-Stamping

(i) Note

This feature is supported on ConnectX-6 Dx and above adapter cards only.

Transmitted packet time-stamping accuracy can be improved when using a timestamp generated at the port level instead of a timestamp generated upon CQE creation. Tx port time-stamping better reflects the actual time of a packet's transmission.

Normal Send queues (SQs) are open with CQE time-stamp support. When this feature is enabled, the driver is expected to open extra Tx port time-stamped SQ per traffic class (TC).

The stream must meet the following conditions in order to be transmitted through a Tx port time-stamped SQ.

- 1. SKBTX_HW_TSTAMP flag was set at tx_flag (SO_TIMESTAMPING was set via setsockopt() or similarly)
- 2. Packet type is:
 - 1. Non-IP, with EtherType of PTP over IEEE 802.3 (0x88f7)

or

2. UDP over IPv4/IPv6

This feature is disabled by default in order to avoid extra SQ memory allocations. The feature can be enabled or disabled using the following command.

ethtool --set-priv-flags <ifs-name> tx_port_ts on / off

PTP Cyc2time Hardware Translation Offload

(i) Note

This feature is supported on ConnectX-6 Dx and above adapter cards only.

Device timestamp can be in one of two modes: real time or free running internal time.

In free running internal time mode, the device clock is not editable in any way. Driver and/or user space must adjust it to the real-time nanosecond values.

In real time mode, the hardware clock device can be adjusted and can provide timestamps which are already translated into real-time nanoseconds.

Both modes are global per device. Once a mode is set, all clock-related features (such as PPS, CQE TS, PCIe bar, etc) will work with the chosen clock mode only.

Free running internal time is the default mode configured in the hardware. The driver will modify the hardware real time clock based on PTP daemon clock adjustments.

Only physical functions are allowed to modify the hardware real-time clock, so PTP daemon adjustments from VFs will be treated as NOP. In case more than one physical function tries to modify the hardware real-time clock, the device will select one of the functions as its designated clock provider. All other input will also be treated as a NOP. The designated clock provide can be replaced by the device if no new adjustments have been received from the current provider after some period.

Timestamp Format

CQE hardware timestamp format for ConnectX-6 Dx and ConnectX-6 Lx NICs is 64 bit, as follows:

{32bit sec, 32 bit nsec}

Configuration

To enable the feature:

- 1. Set REAL_TIME_CLOCK_ENABLE in NV_CONFIG via mlxconfig
- 2. Restart the driver.

Limitations

- Administrator must restart the driver and perform a FW reset for the configuration to take effect. Otherwise, mismatch between HW and driver timestamp mode might occur.
- Once real time mode is activated on a given device (see configuration section), version 5.3 or newer must run on all device functions. Any older driver running on a device function at this configuration will fail to open any traffic queues (RDMA or ETH), hence becoming dysfunctional.
- In real time mode, all device functions must be PTP-synchronized by a single clock domain—do not use multiple GMs for different functions on the same device.
- Regarding hardware clock ownership, the hardware is configured only from a single elected function; other function settings are ignored by the device. There is no indication as to which function is the hardware-clock's owner. After an internal timeout without modifying the hardware clock, a function losses the hardware-clock's ownership and is open to be grasped by any of the functions.
- All PFs/VFs within the same device must sync to the same 1588 master clock. If multiple masters are used, the device will use a single elected function. This might lead to wrong clock representation by device, wrong 1588 TLVs and hiccups on replacement of elected function.
- This feature is supported on ConnectX-6 Dx and above adapter cards only.

RoCE Time-Stamping

RoCE Time-Stamping allows you to stamp packets when they are sent to the wire/received from the wire. The time-stamp is given in raw hardware cycles but could be easily converted into hardware referenced nanoseconds based time. Additionally, it enables you to query the hardware for the hardware time, thus stamp other application's event and compare time.

Query Capabilities

Time-stamping is available if and only the hardware reports it is capable of reporting it. To verify whether RoCE Time-Stamping is available, run ibv_query_device_ex.

For further information, please see ibv_query_device_ex_manual page.

Creating a Time-Stamping Completion Queue

To get time stamps, a suitable extended Completion Queue (CQ) must be created via a special call to ibv_create_cq_ex verb.

For further information, please see ibv_create_cq_ex manual page.

i Note

Time Stamping in not available when CQE zipping is used.

Querying the Hardware Time

Querying the hardware for time is done via the ibv_query_rt_values_ex verb. For example:

For further information, please see ibv_query_rt_values_ex manual page.

One Pulse Per Second (1PPS)

1PPS is a time synchronization feature that allows the adapter to be able to send or receive 1 pulse per second on a dedicated pin on the adapter card using an SMA connector (SubMiniature version A). Only one pin is supported and could be configured as 1PPS in or 1PPS out.

For further information, refer to <u>HowTo Test 1PPS on NVIDIA Adapters</u> Community post.

Flow Steering

Flow steering is a new model which steers network flows based on flow specifications to specific QPs. Those flows can be either unicast or multicast network flows. In order to maintain flexibility, domains and priorities are used. Flow steering uses a methodology of flow attribute, which is a combination of L2-L4 flow specifications, a destination QP and a priority. Flow steering rules may be inserted either by using ethtool or by using InfiniBand verbs. The verbs abstraction uses different terminology from the flow attribute (ibv_flow_attr), defined by a combination of specifications (struct ibv_flow_spec_*).

Flow Steering Support

All flow steering features are enabled in the supported adapter cards. Flow Steering support in InfiniBand is determined according to the MANAGED_FLOW_STEERING flag.

Flow Domains and Priorities

Flow steering defines the concept of domain and priority. Each domain represents a user agent that can attach a flow. The domains are prioritized. A higher priority domain will always supersede a lower priority domain when their flow specifications overlap. Setting a lower priority value will result in a higher priority.

In addition to the domain, there is a priority within each of the domains. Each domain can have at most 2^12 priorities in accordance with its needs.

The following are the domains at a descending order of priority:

- User Verbs allows a user application QP to be attached to a specified flow when using ibv_create_flow and ibv_destroy_flow verbs
- ibv_create_flow

```
struct ibv_flow *ibv_create_flow(struct ibv_qp *qp, struct
ibv_flow_attr
*flow)
```

Input parameters:

- struct ibv_qp the attached QP.
- struct ibv_flow_attr attaches the QP to the flow specified. The flow contains mandatory control parameters and optional L2, L3 and L4 headers. The

optional headers are detected by setting the size and num_of_specs fields:

struct ibv_flow_attr can be followed by the optional flow headers structs:

struct ibv_flow_spec_eth
struct ibv_flow_spec_ipv4
struct ibv_flow_spec_tcp_udp
struct ibv_flow_spec_ipv6

For further information, please refer to the ibv_create_flow man page.

• ibv_destroy_flow

int ibv_destroy_flow(struct ibv_flow *flow_id)

Input parameters:

ibv_destroy_flow requires struct ibv_low which is the return value of ibv_create_flow in case of success.

Output parameters:

Returns 0 on success, or the value of errno on failure.

For further information, please refer to the ibv_destroy_flow man page.

Ethtool

Ethtool domain is used to attach an RX ring, specifically its QP to a specified flow. Please refer to the most recent ethtool man page for all the ways to specify a flow.

Examples:

• ethtool –U eth5 flow-type ether dst 00:11:22:33:44:55 loc 5 action 2

All packets that contain the above destination MAC address are to be steered into rx-ring 2 (its underlying QP), with priority 5 (within the ethtool domain)

• ethtool –U eth5 flow-type tcp4 src-ip 1.2.3.4 dst-port 8888 loc 5 action 2

All packets that contain the above destination IP address and source port are to be steered into rx- ring 2. When destination MAC is not given, the user's destination MAC is filled automatically.

ethtool -U eth5 flow-type ether dst 00:11:22:33:44:55 vlan 45 m 0xf000 loc 5 action 2

All packets that contain the above destination MAC address and specific VLAN are steered into ring 2. Please pay attention to the VLAN's mask 0xf000. It is required in order to add such a rule.

• ethtool -u eth5

Shows all of ethtool's steering rule

When configuring two rules with the same priority, the second rule will overwrite the first one, so this ethtool interface is effectively a table. Inserting Flow Steering rules in the kernel requires support from both the ethtool in the user space and in kernel (v2.6.28).

Accelerated Receive Flow Steering (aRFS)

Receive Flow Steering (RFS) and Accelerated Receive Flow Steering (aRFS) are kernel features currently available in most distributions. For RFS, packets are forwarded based on the location of the application consuming the packet. aRFS boosts the speed of RFS by adding support for the hardware. By usingaRFS(unlike RFS), the packets are directed to a CPU that is local to the thread running the application.

aRFSis an in-kernel-logic responsible for load balancing between CPUs by attaching flows to CPUs that are used by flow's owner applications. This domain allows the aRFS mechanism to use the flow steering infrastructure to support the aRFS logic by implementing the ndo_rx_flow_steer, which, in turn, calls the underlying flow steering mechanism with the aRFS domain.

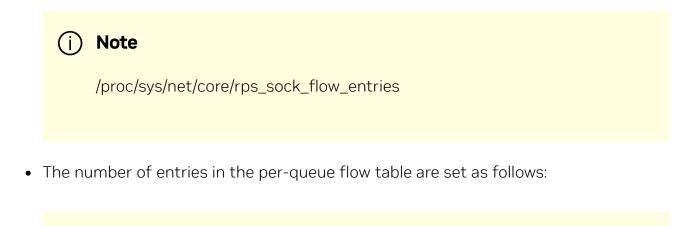
To configure RFS:

≻

Configure the RFS flow table entries (globally and per core).

Note: The functionality remains disabled until explicitly configured (by default it is 0).

• The number of entries in the global flow table is set as follows:



(i) Note /sys/class/net/<dev>/queues/rx-<n>/rps_flow_cnt

Example:

```
# echo 32768 > /proc/sys/net/core/rps_sock_flow_entries
# NUM_CHANNELS=`ethtool -l ens6 | grep "Combined:" | tail -1 | awk
'{print $2}'`
# for f in `seq 0 $((NUM_CHANNELS-1))`; do echo 32768 >
/sys/class/net/ens6/queues/rx-$f/rps_flow_cnt; done
```

To Configure aRFS:

≻

The aRFS feature requires explicit configuration in order to enable it. Enabling the aRFS requires enabling the 'ntuple' flag via the ethtool.

For example, to enable ntuple for ethO, run:

ethtool -K eth0 ntuple on

aRFS requires the kernel to be compiled with the CONFIG_RFS_ACCEL option. This option is available in kernels 2.6.39 and above. Furthermore, aRFS requires Device Managed Flow Steering support.

(i) Note

RFS cannot function if LRO is enabled. LRO can be disabled via ethtool.

Flow Steering Dump Tool

The mlx_fs_dump is a python tool that prints the steering rules in a readable manner. Python v2.7 or above, as well as pip, anytree and termcolor libraries are required to be installed on the host.

Running example:

For further information on the mlx_fs_dump tool, please refer to <u>mlx_fs_dump</u> Community post.

Wake-on-LAN (WoL)

Wake-on-LAN (WoL) is a technology that allows a network professional to remotely power on a computer or to wake it up from sleep mode.

• To enable WoL:

ethtool -s <interface> wol g

• To get WoL:

ethtool <interface> | grep Wake-on Wake-on: g

Where:

"g" is the magic packet activity.

Hardware Accelerated 802.1ad VLAN (Q-in-Q Tunneling)

Q-in-Q tunneling allows the user to create a Layer 2 Ethernet connection between two servers. The user can segregate a different VLAN traffic on a link or bundle different VLANs into a single VLAN. Q-in-Q tunneling adds a service VLAN tag before the user's 802.1Q VLAN tags.

For Q-in-Q support in virtualized environments (SR-IOV), please refer to "<u>Q-in-Q</u> <u>Encapsulation per VF in Linux (VST)</u>".

To enable device support for accelerated 802.1ad VLAN:

1. Turn on the new ethtool private flag "phv-bit" (disabled by default).

\$ ethtool --set-priv-flags eth1 phv-bit on

Enabling this flag sets the phv_en port capability.

2. Change the interface device features by turning on the ethtool device feature "txvlan- stag-hw-insert" (disabled by default).

\$ ethtool -K eth1 tx-vlan-stag-hw-insert on

Once the private flag and the ethtool device feature are set, the device will be ready for 802.1ad VLAN acceleration.



The "phv-bit" private flag setting is available for the Physical Function (PF) only.

The Virtual Function (VF) can use the VLAN acceleration by setting the "tx-vlan-stag-hw-insert" parameter only if the private flag "phv-bit" is enabled by the PF. If the PF enables/disables the "phv-bit" flag after the VF driver is up, the configuration will take place only after the VF driver is restarted.

VLAN Stripping in Linux Verbs

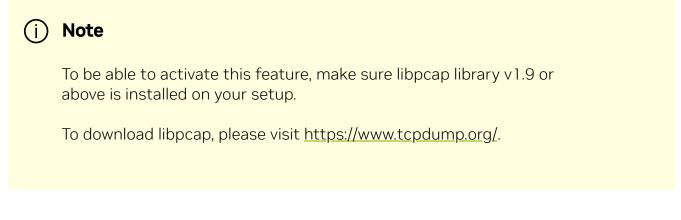
(i) Note

This capability is now accessible from userspace using the verbs.

VLAN stripping adds access to the device's ability to offload the Customer VLAN (cVLAN) header stripping from an incoming packet, thus achieving acceleration of VLAN handing in receive flow.

It is configured per WQ option. You can either enable it upon creation or modify it later using the appropriate verbs (ibv_create_wq/ibv_modify_wq).

Offloaded Traffic Sniffer



Offloaded Traffic Sniffer allows bypass kernel traffic (such as RoCE, VMA, and DPDK) to be captured by existing packet analyzer, such as tcpdump.

To capture the interface's bypass kernel traffic, run tcpdump on the RDMA device.

For examples on how to dump RDMA traffic using the Inbox tcpdump tool for ConnectX-4 adapter cards and above, click <u>here</u>.

j) Note

Note that enabling Offloaded Traffic Sniffer can cause bypass kernel traffic speed degradation.

) Note

In case you do not wish to install libpcap on your setup, you can use docker to run the tcpdump. For further information, please see <u>https://hub.docker.com/r/mellanox/tcpdump-rdma</u>.

Dump Configuration

This feature helps dumping driver and firmware configuration using ethtool. It creates a backup of the configuration files into a specified dump file.

Dump Parameters (Bitmap Flag)

The following bitmap parameters are used to set the type of dump. If a value is not set, the default value used is "0".

Bitmap Parameters

Value	Description
1	MST dump
2 Ring dump (Software context information for SQs, EQs, RQs, CQs)	
3 MST dump + Ring dump (1+2)	
4	Clear this parameter

Configuration

In order to configure this feature, follow the steps below:

1. Set the dump bitmap parameter by running -W (uppercase) with the desired bitmap parameter value (see Bitmap Parameters table above). In the following example, the bitmap parameter value is 3.

ethtool -W ens1f0 3

2. Dump the file by running -w (lowercase) with the desired configuration file name.

ethtool -w ens1f0 data /tmp/dump.bin

3. [**Optional**] To get the bitmap parameter value, version and size of the dump, run the command above without the file name.

```
ethtool -w ens1f0
flag: 3, version: 1, length: 4312
```

4. To open the dump file, run:

```
mlnx_dump_parser -f /tmp/dump.bin -m mst_dump_demo.txt -r
ring_dump_demo.txt
Version: 1 Flag: 3 Number of blocks: 123 Length 327584
MCION module number: 0 status: | present |
DRIVER VERSION: 1-23 (03 Mar 2015)
DEVICE NAME 0000:81:00.0:ens1f0
Parsing Complete!
```

where:

-f	For the file to be parsed (the file that was just created)
-m	For the mst dump file
-r	For the ring dump file

For further information, refer to <u>HowTo Dump Driver Configuration (via ethtool)</u> Community post.

Output:

```
# mlnx_dump_parser -f /tmp/dump.bin -m mst_dump_demo.txt -r
ring_dump_demo.txt
Version: 1 Flag: 3 Number of blocks: 123 Length 327584
MCION module number: 0 status: | present |
DRIVER VERSION: 1-23 (03 Mar 2015)
DEVICE NAME 0000:81:00.0:ens1f0
Parsing Complete!
```

- 5. Open the files.
 - 1. The MST dump file will look as follows. In order to analyze it, contact <u>NVIDIA</u> <u>Support</u>.

```
cat mst_dump_demo.txt
0x00000000 0x01002000
0x0000004 0x00000000
0x0000000c 0x00000000
0x00000010 0x00000000
0x00000014 0x0000000
0x00000018 0x0000000
....
```

2. The Ring dump file can help developers debug ring-related issues, and it looks as follows:

cat ring_dump_demo.txt SQ TYPE: 3, WQN: 102, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024... SQ TYPE: 3, WQN: 102, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 65536, GROUP_IP: 0 CQ TYPE: 5, WQN: 20, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 1024, GROUP_IP: 0 RQ TYPE: 4, WQN: 103, PI: 15, CI: 0, STRIDE: 5, SIZE: 16, WQE_NUM: 512, GROUP_IP: 0 CQ TYPE: 5, WQN: 21, PI: 0, CI: 0, STRIDE: 6, SIZE: 16384, WQE_NUM: 16384, GROUP_IP: 0 EQ TYPE: 6, CI: 1, SIZE: 0, IRQN: 109, EQN: 19, NENT: 2048, MASK: 0, INDEX: 0, GROUP_ID: 0 SQ TYPE: 3, WQN: 106, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 65536, GROUP_IP: 1 CQ TYPE: 5, WQN: 23, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 1024, GROUP_IP: 1 RQ TYPE: 4, WQN: 107, PI: 15, CI: 0, STRIDE: 5, SIZE: 16, WQE_NUM: 512, GROUP_IP: 1

CQ TYPE: 5, WQN: 24, PI: 0, CI: 0, STRIDE: 6, SIZE: 16384, WQE_NUM: 16384, GROUP_IP: 1 EQ TYPE: 6, CI: 1, SIZE: 0, IRQN: 110, EQN: 20, NENT: 2048, MASK: 0, INDEX: 1, GROUP_ID: 1 SQ TYPE: 3, WQN: 110, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 65536, GROUP_IP: 2 CQ TYPE: 5, WQN: 26, PI: 0, CI: 0, STRIDE: 6, SIZE: 1024, WQE_NUM: 1024, GROUP_IP: 2 RQ TYPE: 4, WQN: 111, PI: 15, CI: 0, STRIDE: 5, SIZE: 16, WQE_NUM: 512, GROUP_IP: 2 CQ TYPE: 5, WQN: 27, PI: 0, CI: 0, STRIDE: 6, SIZE: 16384, WQE_NUM: 16384, GROUP_IP: 2 ...

Local Loopback Disable

Local Loopback Disable feature allows users to force the disablement of local loopback on the virtual port (vport). This disables both unicast and multicast loopback in the hardware.

To enable Local Loopback Disable, run the following command:

echo 1 > /sys/class/net/<ifname>/settings/force_local_lb_disable"

To disable Local Loopback Disable, run the following command:

echo 0 > /sys/class/net/<ifname>/settings/force_local_lb_disable"

(i) Note

Kernel Transport Layer Security (kTLS) Offloads



This feature is supported on ConnectX-6 Dx crypto cards only.

Overview

Transport Layer Security (TLS) is a widely-deployed protocol used for securing TCP connections on the Internet. TLS is also a required feature for HTTP/2, the latest web standard. Kernel implementation of TLS (kTLS) provides new opportunities for offloading the protocol into the hardware.

TLS data-path offload allows the NIC to accelerate encryption, decryption and authentication of AES-GCM. TLS offload handles data as it goes through the device without storing any data, but only updating context. If the packet cannot be encrypted/decrypted by the device, then a software fallback handles the packet.

Establishing a kTLS Connection

To avoid unnecessary complexity in the kernel, the TLS handshake is kept in the user space. A full TLS connection using the socket is done using the following scheme:

- 1. Call connect() or accept() on a standard TCP file descriptor.
- 2. Use a user space TLS library to complete a handshake.
- 3. Create a new KTLS socket file descriptor.

- 4. Extract the TLS Initialization Vectors (IVs), session keys, and sequence IDs from the TLS library. Use the setsockopt function on the kTLS file descriptor (FD) to pass them to the kernel.
- 5. Use standard read(), write(), sendfile() and splice() system calls on the KTLS FD.

Drivers can offer Tx and Rx packet encryption/decryption offload from the kernel into the NIC hardware. Upon receipt of a non-data TLS message (a control message), the kTLS socket returns an error, and the message is left on the original TCP socket instead. The kTLS socket is automatically unattached. Transfer of control back to the original encrypted FD is done by calling getsockopt to receive the current sequence numbers, and inserting them into the TLS library.

Kernel Support

For support in the kernel, make sure the following flags are set as follows.

- CONFIG_TLS=y
- CONFIG_TLS_DEVICE=y | m

j) Note

For kTLS **Tx** device offloads with OFED drivers, kernel TLS module (kernel/net/tls) must be aligned to kernel v5.3 and above.

For kTLS ${\bf Rx}$ device offloads with OFED drivers, kernel TLS module (kernel/net/tls) must be aligned to kernel v5.9 and above.

Configuring kTLS Offloads

To enable kTLS Tx offload, run:

ethtool -K <ifs> tls-hw-tx-offload on

ethtool -K <ifs> tls-hw-rx-offload on

For further information on TLS offloads, please visit the following kernel documentation:

- <u>https://www.kernel.org/doc/html/latest/networking/tls-offload.html</u>
- <u>https://www.kernel.org/doc/html/latest/networking/tls.html#kernel-tls</u>

OpenSSL with kTLS Offload

OpenSSL version 3.0.0 or above is required to support kTLS TX/RX offloads.

Supported OpenSSL version is available to download from distro packages, or can be downloaded and compiled from the OpenSSL github.

IPsec Crypto Offload

i Note

This feature is supported on crypto-enabled products of BlueField-2 DPUs, and ConnectX-6 Dx and ConnectX-7 adapters (but not of ConnectX-6 or ConnectX-6 Lx).

Newer/future crypto-enabled DPU and adapter product generations should also support the feature, unless explicitly stated in their documentation.

i) Note

For NVIDIA BlueField-2 DPUs and ConnectX-6 Dx adapters Only: If your target application will utilize bandwidth of 100Gb/s or higher, where a substantial part of the bandwidth will be allocated for IPsec traffic, please refer to the NVIDIA BlueField-2 DPUs Product Release Notes or NVIDIA ConnectX-6 Dx Adapters Product Release Notes document to learn about a potential bandwidth limitation. To access the relevant product release notes, please contact your NVIDIA sales representative.

Overview and Configuration

IPsec crypto offload feature, also known as IPsec inline offload or IPsec aware offload feature enables the user to offload IPsec crypto encryption and decryption operations to the hardware.

Note that the hardware implementation only supports AES-GCM encryption scheme.

To enable the feature, support in both kernel and adapter firmware is required.

• For support in the kernel, make sure the following flags are set as follows.

CONFIG_XFRM_OFFLOAD=y CONFIG_INET_ESP_OFFLOAD=m CONFIG_INET6_ESP_OFFLOAD=m

Note: These flags are enabled by default in RedHat 8 and Ubuntu 18.04.

• For support in the firmware, make sure the below string is found in the dmesg.

mlx5e: IPSec ESP acceleration enabled

Configuring Security Associations for IPsec Offloads

To program the inline offload security associations (SA), add the option "offload dev <netdev interface> dir out/in" in the "ip xfrm state" command for transmitting and receiving SA.

Transmit inline offload SA xfrm command example:

sudo ip xfrm state add src 192.168.1.64/24 dst 192.168.1.65/24 proto esp spi 0x46dc6204 reqid 0x46dc6204 mode transport aead 'rfc4106(gcm(aes))' 0x60bd6c3eafba371a46411830fd56c53af93883261ed1fb26767820ff493f43ba35b0dcca offload dev p4p1 dir out sel src 192.168.1.64 dst 192.168.1.65

Receive inline offload SA xfrm command example:

sudo ip xfrm state add src 192.168.1.65/24 dst 192.168.1.64/24 proto esp spi 0xaea0846c reqid 0xaea0846c mode transport aead 'rfc4106(gcm(aes))' 0x81d5c3167c912c1dd50dab0cb4b6d815b6ace8844304db362215a258cd19deda8f89de offload dev p4p1 dir in sel src 192.168.1.65 dst 192.168.1.64

Setting xfrm Policies Example

First server:

+ sudo ip xfrm state add src 192.168.1.64/24 dst 192.168.1.65/24 proto esp spi 0x28f39549 reqid 0x28f39549 mode transport aead 'rfc4106(gcm(aes))' 0x492e8ffe718a95a00c1893ea61afc64997f4732848ccfe6ea07db483175cb18de9ae411a offload dev enp4s0 dir out sel src 192.168.1.64 dst 192.168.1.65 + sudo ip xfrm state add src 192.168.1.65/24 dst 192.168.1.64/24 proto esp spi 0x622a73b4 reqid 0x622a73b4 mode transport aead 'rfc4106(gcm(aes))' 0x093bfee2212802d626716815f862da31bcc7d9c44cfe3ab8049e7604b2feb1254869d25 offload dev enp4s0 dir in sel src 192.168.1.65 dst 192.168.1.64 + sudo ip xfrm policy add src 192.168.1.64 dst 192.168.1.65 dir out tmpl src 192.168.1.64/24 dst 192.168.1.65/24 proto esp reqid 0x28f39549 mode transport + sudo ip xfrm policy add src 192.168.1.65 dst 192.168.1.64 dir in tmpl src 192.168.1.65/24 dst 192.168.1.65/24 proto esp reqid 0x622a73b4 mode transport + sudo ip xfrm policy add src 192.168.1.65 dst 192.168.1.64 dir fwd tmpl src 192.168.1.65/24 dst 192.168.1.64/24 proto esp reqid 0x622a73b4 mode transport

Second server:

```
+ ssh -A -t root@l-csi-0921d /bin/bash
+ set -e
+ '[' 0 == 1 ']'
+ sudo ip xfrm state add src 192.168.1.64/24 dst 192.168.1.65/24 proto
esp spi 0x28f39549 reqid 0x28f39549 mode transport aead
'rfc4106(gcm(aes))' 0x492e8ffe718a95a00c1893ea61afc64997f4732848ccfe6ea07db483175cb18de9ae411a
offload dev enp4s0 dir in sel src 192.168.1.64 dst 192.168.1.65
+ sudo ip xfrm state add src 192.168.1.65/24 dst 192.168.1.64/24 proto
esp spi 0x622a73b4 regid 0x622a73b4 mode transport aead
'rfc4106(gcm(aes))' 0x093bfee2212802d626716815f862da31bcc7d9c44cfe3ab8049e7604b2feb1254869d25
offload dev enp4s0 dir out sel src 192.168.1.65 dst 192.168.1.64
+ sudo ip xfrm policy add src 192.168.1.65 dst 192.168.1.64 dir out
tmpl src 192.168.1.65/24 dst 192.168.1.64/24 proto esp regid 0x622a73b4
mode transport
+ sudo ip xfrm policy add src 192.168.1.64 dst 192.168.1.65 dir in tmpl
src 192.168.1.64/24 dst 192.168.1.65/24 proto esp regid 0x28f39549 mode
transport
+ sudo ip xfrm policy add src 192.168.1.64 dst 192.168.1.65 dir fwd
tmpl src 192.168.1.64/24 dst 192.168.1.65/24 proto esp regid 0x28f39549
mode transport
+ echo 'IPSec tunnel configured successfully'
```

IPsec Full Offload

(i) Note

This feature is supported on crypto-enabled products of BlueField-2 DPUs, as well as on ConnectX-6 Dx, ConnectX-6 Lx and ConnectX-7 adapter cards. Note that it is not supported on ConnectX-6 cards.

Newer/future crypto-enabled DPU and adapter product generations should also support this feature, unless explicitly stated otherwise in their documentation.

(i) Note

When using NVIDIA® BlueField®-2 DPUs and NVIDIA® ConnectX®-6 Dx adapters only: If your target application utilizes 100Gb/s or a higher bandwidth, where a substantial part of the bandwidth is allocated for IPsec traffic, please refer to the relevant DPU or adapter card Product Release Notes to learn about a potential bandwidth limitation. To access the Release Notes, visit <u>https://docs.nvidia.com/networking/</u>, or contact your NVIDIA sales representative.

(i) Note

ConnectX-6 Dx adapters only support Full Offload: Encrypted Overlay (where a Hypervisor controls IPsec offload - See for example OVS IPsec - <u>https://docs.openvswitch.org/en/latest/tutorials/ipsec/</u>) in a Linux OS with NVIDIA drivers.

i) Note

This feature requires Linux kernel v6.6, or higher.

This feature is designed to enable IPsec full offload in switchdev mode. The ip-xfrm command is used to configure IPsec states and policies, and it is similar to legacy mode configuration. However, there are several limitations to the use of full offload in this mode:

- 1. Only IPsec Transport Mode and Tunnel Mode are supported.
- 2. The first IPsec TX state/policy is not allowed to be offloaded if any offloaded TC rule exists, and the same applies for the first RX state/policy. More specifically, IPsec RX/TX tables must be created before offloading any TC rule. For this reason, it is a common practice to configure IPsec rules before adding any TC rule.

Following is an example for IPsec configuration with a VXLan tunnel:

• Enable switchdev mode:

```
echo 1 > /sys/class/net/$PF0 /device/sriov_numvfs
echo 0000:08:00.2 > /sys/bus/pci/drivers/mlx5_core/unbind
devlink dev param set pci/0000:08:00.0 name flow_steering_mode
value dmfs cmode runtime
devlink dev eswitch set pci/0000:08:00.0 mode switchdev
echo 0000:08:00.2 > /sys/bus/pci/drivers/mlx5_core/bind
```

• Configure PF/VF/REP netdevices, and place a VF in a namespace:

```
ifconfig $PF $LOCAL_TUN/16 up
ip l set dev $PF mtu 2000
ifconfig $REP up
ip netns add ns0
ip link set dev $VF netns ns0
ip netns exec ns0 ifconfig $VF $IP/16 up
```

• Configure IPsec states and policies:

```
ip xfrm state add src $LOCAL_TUN/16 dst $REMOTE_IP/16 proto
esp spi 0xb29ed314 regid 0xb29ed314 mode transport aead
'rfc4106(gcm(aes))' 0x20f01f80a26f633d85617465686c32552c92c42f 128 offload
packet dev $PF dir out sel src $LOCAL_TUN/16 dst
$REMOTE_IP/16 flag esn replay-window 64
ip xfrm state add src $REMOTE_IP/16 dst $LOCAL_TUN/16 proto
esp spi 0xc35aa26e regid 0xc35aa26e mode transport aead
'rfc4106(gcm(aes))' 0x6cb228189b4c6e82e66e46920a2cde39187de4ba 128 offload
packet dev $PF dir in sel src $REMOTE_IP/16 dst $LOCAL_TUN/16
flag esn replay-window 64
ip xfrm policy add src $LOCAL_TUN dst $REMOTE_IP offload
packet dev $PF dir out tmpl src $LOCAL_TUN/16 dst
$REMOTE_IP/16 proto esp regid 0xb29ed314 mode transport
priority 12
ip xfrm policy add src $REMOTE_IP dst $LOCAL_TUN offload
packet dev $PF dir in tmpl src $REMOTE_IP/16 dst
$LOCAL_TUN/16 proto esp regid 0xc35aa26e mode transport priority
12
```

• Configure Openvswitch:

```
ovs-vsctl add-br br-ovs
ovs-vsctl add-port br-ovs $REP
ovs-vsctl add-port br-ovs vxlan1 -- set interface vxlan1
type=vxlan options:local_ip=$LOCAL_TUN
options:remote_ip=$REMOTE_IP options:key=$VXLAN_ID
options:dst_port=4789
```

IPsec Full Offload for RDMA Traffic

This IPsec Full Offload for RDMA Traffic option provides a significant performance improvement compared to the software IPsec counterpart, and enables the use of IPsec over RoCE packets, which are outside the network stack and cannot be used without full hardware offload. As a result, users can leverage the benefits of the IPsec protocol with RoCE V2, even when using SR-IOV VFs.

The configuration steps for this feature should be identical to the steps mentioned above, but if this feature is supported, the traffic that will be sent can also be RoCEV2 IPsec traffic.

To configure this feature:

- 1. Configure an SR-IOV VF normally, and add its OVS/TC rules.
- 2. Enable IPsec over VF. For more information, please see IPsec Functionality.
- 3. Configure IPsec policies and states on the relevant VF net device. This should be identical to the software configuration of IPsec rules, which can be done using one of the following implementation options:

Command	Offload Request Parameter	
iproute2 ip xfrm	offload packet	
libreswan		
strongswan	nic-offload=packet	

j) Note

For this feature to work, switchdev mode and dmfs steering mode must be enabled.

• The following is a full minimalistic configuration example using iproute, whereas PFO is the netdevice PF, FO_REP is the VF representor, and NIC is the VF netdevice to configure IPsec over:

1. echo 1 > /sys/class/net/\$PF0 /device/sriov_numvfs

2. echo 0000:08:00.2 >

/sys/bus/pci/drivers/mlx5_core/unbind

devlink dev eswitch set pci/0000:08:00.0 mode switchdev
devlink dev param set pci/0000:08:00.0 name
flow_steering_mode value dmfs cmode runtime

5. devlink port function set pci/0000:08:00.0/1
ipsec_packet enable

6. echo 0000:08:00.2 > /sys/bus/pci/drivers/mlx5_core/bind

7. tc qdisc add dev \$PF0 ingress

tc qdisc add dev \$VF0_REP ingress

tc filter add dev \$PF0 parent ffff: protocol 802.1q chain 0
flower vlan_id 10 vlan_ethtype 802.1q cvlan_id 5 action vlan
pop action vlan pop action mirred egress redirect dev
\$VF0_REP

tc filter add dev \$VF0_REP parent ffff: protocol all chain 0 flower action vlan push protocol 802.1q id 5 action vlan push protocol 802.1q id 10 action mirred egress redirect dev \$PF0

8. ifconfig \$PF0 \$PF_IP/24 up
ifconfig \$NIC \$LOC_IP/\$SUB_NET up
ip link set dev \$VF_REP up
9. ip xfrm state flush
ip xfrm policy flush

• Configure ipsec states and policies:

#states

ip -4 xfrm state add src \$LOC_IP/\$SUB_NET dst \$REMOTE_IP/\$SUB_NET proto esp spi 1000 reqid 10000 aead 'rfc4106(gcm(aes))' 0x010203047aeaca3f87d060a12f4a4487d5a5c335 128 mode transport sel src \$LOC_IP dst \$REMOTE_IP offload packet dev \$NIC dir out ip -4 xfrm state add src \$REMOTE_IP/\$SUB_NET dst \$LOC_IP/\$SUB_NET proto esp spi 1001 reqid 10001 aead 'rfc4106(gcm(aes))' 0x010203047aeaca3f87d060a12f4a4487d5a5c335 128 mode transport sel src \$REMOTE_IP dst \$LOC_IP offload packet dev \$NIC dir in #policies ip -4 xfrm policy add src \$LOC_IP dst \$REMOTE_IP offload packet dev \$NIC dir out tmpl src \$LOC_IP/\$SUB_NET dst \$REMOTE_IP/\$SUB_NET proto esp reqid 10000 mode transport ip -4 xfrm policy add src \$REMOTE_IP dst \$LOC_IP offload packet dev \$NIC dir in tmpl src \$REMOTE_IP dst \$LOC_IP offload packet dev \$NIC dir in tmpl src \$REMOTE_IP/\$SUB_NET dst \$LOC_IP/\$SUB_NET proto esp reqid 10001 mode transport ip -4 xfrm policy add src \$REMOTE_IP dst \$LOC_IP dir fwd tmpl src \$REMOTE_IP/\$SUB_NET dst \$LOC_IP dir fwd tmpl src \$REMOTE_IP/\$SUB_NET dst \$LOC_IP/\$SUB_NET proto esp reqid 10001 mode transport

Note that the configuration above is for one side only, yet IPsec must be configured for both sides in order for them to communicate properly. The configuration for the other side should be almost identical, but Step 9 would be configured in an asymmetrical way, meaning the first policy would look the following, and all other states/policies would be adjusted accordingly:

ip -4 xfrm state add src \$LOC_IP/\$SUB_NET dst \$REMOTE_IP/\$SUB_NET
proto esp spi 1001 reqid 10001 aead
'rfc4106(gcm(aes))' 0x010203047aeaca3f87d060a12f4a4487d5a5c335 128 mode transport
sel src \$LOC_IP dst \$REMOTE_IP offload packet dev \$NIC dir out

Once this step is completed, you can send any RoCE traffic of your choice between the two machines with configured IPsec. For example, ibv_rc_pingpong -g 3 -d VF_device : on one side, and ibv_rc_pingpong -g 3 -d VF_device \$IP_OF_OTHER_SIDE : on the other side.

Finally, you can verify that the traffic was encrypted using IPsec by using the ipsec counters:

MACsec Full Offload

MACsec Full offload feature, also known as MACsec inline Full offload, enables the user to offload MACsec crypto encryption and decryption, MACsec headers encapsulation and decapsulation, and Anti replay operations to the hardware.

j) Note

Hardware implementation supports GCM-AES & GCM-AES-XPN encryption schemes and is supported with ConnectX-7 onwards.

i) Note

MACsec introduced in MOFED v5.9 requires a minimal Kernel version of 6.1.

To enable the feature, support in both kernel and adapter firmware is required.

For support in the kernel, make sure the following flags are set as follows:

- CONFIG_MACSEC=y
- CONFIG_MLX5_EN_MACSEC=y

For support in firmware use the following version:

• xx.34.0364 and up

Configurations

IProute2 Configuration

Configuring Physical Interface

Client side:

- ip address flush <physical_device>
- ip address add <client_physical_device_ip> dev <physical interface>
- ip link set dev <physical_device>up

Server side:

- ip address flush <physical_device>
- ip address add <server_physical_device_ip> dev <physical interface>
- ip link set dev <physical_device>up

Add MACsec Device

Client side:

ip link add link <physical_device> <macsec_device> type macsec sci <client_sci> client on

Server side:

ip link add link <physical_device> <macsec_device> type macsec sci <server_sci> client on

Offload MACsec Device

Client side:

• ip macsec offload <macsec_device> mac

Server side:

• ip macsec offload <macsec_device> mac

Add MACsec rules:

Client side:

- ip macsec add <macsec_device> tx sa <sa_num>pn <inital_packet_number>on key
 <client_key_id> <client_key>
- ip macsec add <macsec_device> rx sci <server_sci> on
- ip macsec add <macsec_device> rx sci <server_sci>sa <sa_num> pn
 <inital_packet_number> on key <server_key_id> <server_key>

Server side:

- ip macsec add <macsec_device> tx sa <sa_num>pn <inital_packet_number>on key <server_key_id> <server_key>
- ip macsec add <macsec_device> rx sci <client_sci> on
- i p macsec add <macsec_device> rx sci <client_sci>sa <sa_num> pn
 <inital_packet_number> on key <client_key_id> <client_key>

Configure MACsec Device IPs:

Client side:

- ip address flush <macsec_device>
- ip address add <client_macsec_device_ip> dev <macsec_device>
- ip link set dev <macsec_device> up

Server side:

- ip address flush <macsec_device>
- ip address add <server_macsec_device_ip> dev <macsec_device>
- ip link set dev <macsec_device> up

Configuration Example

Client side:

- ip address flush enp8s0f0
- ip address add 1.1.1.1/24 dev enp8s0f0
- ip link set dev enp8s0f0 up
- ip link add link enp8s0f0 macsec0 type macsec sci 1 encrypt on
- ip macsec offload macsec0 mac
- ip macsec add macsec0 tx sa 0 pn 1 on key 00 dffafc8d7b9a43d5b9a3dfbbf6a30c16
- ip macsec add macsec0 rx sci 2 on
- ip macsec add macsec0 rx sci 2 sa 0 pn 1 on key 00 ead3664f508eb06c40ac7104cdae4ce5
- ip address flush macsec0
- ip address add 2.2.2.1/24 dev macsec0
- ip link set dev macsec0 up

Server side:

- ip link del macsec0
- ip address flush enp8s0f0
- ip address add 1.1.1.2/24 dev enp8s0f0
- ip link set dev enp8s0f0 up
- ip link add link enp8s0f0 macsec0 type macsec sci 2 encrypt on
- ip macsec offload macsec0 mac
- ip macsec add macsec0 tx sa 0 pn 1 on key 00 ead3664f508eb06c40ac7104cdae4ce5

- ip macsec add macsec0 rx sci 1 on
- ip macsec add macsec0 rx sci 1 sa 0 pn 1 on key 00 dffafc8d7b9a43d5b9a3dfbbf6a30c16
- ip address flush macsec0
- ip address add 2.2.2.2/24 dev macsec0
- ip link set dev macsec0 up

(i) Note

- Use: "ip macsec show" command to check configuration
- To make sure traffic is offloaded, check MACsec counters: "ethtool -S <physical_device> | grep macsec"

Additional Resources

Linux Manual page: linux_manual

InfiniBand Network

The chapter contains the following sections:

- InfiniBand Interface
- <u>NVIDIA SM</u>
- <u>QoS Quality of Service</u>
- IP over InfiniBand (IPoIB)

- Advanced Transport
- Optimized Memory Access
- <u>NVIDIA PeerDirect</u>
- <u>CPU Overhead Distribution</u>
- Out-of-Order (OOO) Data Placement
- IB Router
- MAD Congestion Control

InfiniBand Interface

Port Type Management

For information on port type management of ConnectX-4 and above adapter cards, please refer to <u>Port Type Management/VPI Cards Configuration</u> section.

RDMA Counters

- RDMA counters are available only through sysfs located under:
 - #/sys/class/infiniband/<device>/ports/*/hw_counters/
 - #/sys/class/infiniband/<device>/ports/*/counters

For mlx5 port and RDMA counters, refer to the <u>Understanding mlx5 Linux Counters</u> Community post.

NVIDIA SM

NVIDIA SM is an InfiniBand compliant Subnet Manager (SM). It is provided as a fixed flow executable called "<u>opensm</u>", accompanied by a testing application called osmtest. NVIDIA SM implements an InfiniBand compliant SM according to the InfiniBand Architecture Specification chapters: Management Model, Subnet Management, and Subnet Administration.

OpenSM Application

OpenSM is an InfiniBand compliant Subnet Manager and Subnet Administrator that runs on top of the NVIDIA OFED stack. OpenSM performs the InfiniBand specification's required tasks for initializing InfiniBand hardware. One SM must be running for each InfiniBand subnet.

OpenSM defaults were designed to meet the common case usage on clusters with up to a few hundred nodes. Thus, in this default mode, OpenSM will scan the IB fabric, initialize it, and sweep occasionally for changes.

OpenSM attaches to a specific IB port on the local machine and configures only the fabric connected to it. (If the local machine has other IB ports, OpenSM will ignore the fabrics connected to those other ports). If no port is specified, opensm will select the first "best" available port. opensm can also present the available ports and prompt for a port number to attach to.

By default, the OpenSM run is logged to/var/log/opensm.log. All errors reported in this log file should be treated as indicators of IB fabric health issues. (Note that when a fatal and non-recoverable error occurs, OpenSM will exit). opensm.log should include the message "SUBNET UP" if OpenSM was able to set up the subnet correctly.

Syntax

opensm [OPTIONS]

For the complete list of OpenSM options, please run:

```
opensm --help / -h / -?
```

Environment Variables

The following environment variables control OpenSM behavior:

• OSM_TMP_DIR - controls the directory in which the temporary files generated by OpenSM are created. These files are: opensm-subnet.lst, opensm.fdbs, and opensm.mcfdbs. By default, this directory is /var/log.

• OSM_CACHE_DIR - opensm stores certain data to the disk such that subsequent runs are consistent. The default directory used is /var/cache/opensm. The following file is included in it:

guid2lid - stores the LID range assigned to each GUID

Signaling

When OpenSM receives a HUP signal, it starts a new heavy sweep as if a trap has been received or a topology change has been found.

Also, SIGUSR1 can be used to trigger a reopen of /var/log/opensm.log for logrotate purposes.

Running OpenSM as Daemon

OpenSM can also run as daemon. To run OpenSM in this mode, enter:

```
host1# service opensmd start
```

osmtest

osmtest is a test program for validating the InfiniBand Subnet Manager and Subnet Administrator. osmtest provides a test suite for opensm. It can create an inventory file of all available nodes, ports, and PathRecords, including all their fields. It can also verify the existing inventory with all the object fields and matches it to a pre-saved one.

osmtest has the following test flows:

- Multicast Compliancy test
- Event Forwarding test
- Service Record registration test

- RMPP stress test
- Small SA Queries stress test

For further information, please refer to the tool's man page.

Partitions

OpenSM enables the configuration of partitions (PKeys) in an InfiniBand fabric. By default, OpenSM searches for the partitions configuration file under the name /etc/opensm/partitions.conf. To change this filename, you can use opensm with the '--Pconfig' or '-P' flags.

The default partition is created by OpenSM unconditionally, even when a partition configuration file does not exist or cannot be accessed.

The default partition has a P_Key value of 0x7fff. The port out of which runs OpenSM is assigned full membership in the default partition. All other end-ports are assigned partial membership.

(i) Note

- Adding a new partition to the partition.conf file, does not require SM restart, but signalling SM process via a HUP signal (e.g pkill -HUP opensm).
- The default partition cannot be removed.

i) Note

Adjustments to the Port GUIDs, including additions, removals, or membership alterations (denoted as "<PortGUID>=[full|limited|both]" in the "Partition Definition") can be applied with a HUP signal to the Subnet Manager process (e.g pkill -HUP opensm).

Marning

Performing changes in the ipoib_bc_flags (ipoib/sl/scope/rate/mtu) and mgroup flags of an existing partition requires a restart of the Subnet Manager to take effect.

File Format

i) Note

Line content followed after '#' character is comment and ignored by parser.

General File Format

<Partition Definition>:\[<newline>\]<Partition Properties>

• <Partition Definition>:

[PartitionName][=PKey][,indx0][,ipoib_bc_flags]
[,defmember=full|limited]

where:

PartitionName	String, will be used with logging. When omitted empty string will be used.
РКеу	P_Key value for this partition. Only low 15 bits will be used. When omitted will be auto-generated.
indx0	Indicates that this pkey should be inserted in block 0 index 0.

ipoib_bc_flags	Used to indicate/specify IPoIB capability of this partition.
defmember=full limit ed both	Specifies default membership for port GUID list. Default is limited.

ipoib_bc_flags are:

ipoib	Indicates that this partition may be used for IPoIB, as a result the IPoIB broadcast group will be created with the flags given, if any.	
rate= <val></val>	Specifies rate for this IPoIB MC group (default is 3 (10GBps))	
mtu= <val></val>	Specifies MTU for this IPoIB MC group (default is 4 (2048))	
s = <val></val>	Specifies SL for this IPoIB MC group (default is 0)	
scope= <val></val>		

• <Partition Properties>:

```
\[<Port list>|<MCast Group>\]* | <Port list>
```

• <Port List>:

<Port Specifier>[,<Port Specifier>]

• <Port Specifier>:

where

PortG	G GUID of partition member EndPort. Hexadecimal numbers should start from 0x	
UID	decimal numbers are accepted too.	

full,	Indicates full and/or limited membership for this both port. When omitted (or	
limite	unrecognized) limited membership is assumed. Both indicate full and limited	
d	membership for this port.	

• <MCast Group>:

mgid=gid[,mgroup_flag]*<newline>

where:

mgid=gid	gid specified is verified to be a Multicast address IP groups are verified to match the rate and mtu of the broadcast group. The P_Key bits of the mgid for IP groups are verified to either match the P_Key specified in by "Partition Definition" or if they are 0x0000 the P_Key will be copied into those bits.	
mgroup_fla g	rate= <val></val>	Specifies rate for this MC group (default is 3 (10GBps))
	mtu= <val></val>	Specifies MTU for this MC group (default is 4 (2048))
	sl= <val></val>	Specifies SL for this MC group (default is 0)
	scope= <val></val>	Specifies scope for this MC group (default is 2 (link local)). Multiple scope settings are permitted for a partition. NOTE: This overwrites the scope nibble of the specified mgid. Furthermore specifying multiple scope settings will result in multiple MC groups being created.
	qkey= <val></val>	Specifies the Q_Key for this MC group (default: 0x0b1b for IP groups, 0 for other groups)
	tclass= <val></val>	Specifies tclass for this MC group (default is 0)
	FlowLabel = <val></val>	Specifies FlowLabel for this MC group (default is 0)

Note that values for rate, MTU, and scope should be specified as defined in the IBTA specification (for example, mtu=4 for 2048). To use 4K MTU, edit that entry to "mtu=5" (5 indicates 4K MTU to that specific partition).

PortGUIDs list:

PortGUID GUID of partition member EndPort. Hexadecimal numbers should start from 0x, decimal numbers are accepted too. full or limited indicates full or limited membership for this port. When omitted (or unrecognized) limited membership is assumed.

There are some useful keywords for PortGUID definition:

- 'ALL_CAS' means all Channel Adapter end ports in this subnet
- 'ALL_VCAS' means all virtual end ports in the subnet
- 'ALL_SWITCHES' means all Switch end ports in this subnet
- 'ALL_ROUTERS' means all Router end ports in this subnet
- 'SELF' means subnet manager's port. An empty list means that there are no ports in this partition

Notes:

- White space is permitted between delimiters ('=', ',';';').
- PartitionName does not need to be unique, PKey does need to be unique. If PKey is repeated then those partition configurations will be merged and the first PartitionName will be used (see the next note).
- It is possible to split partition configuration in more than one definition, but then PKey should be explicitly specified (otherwise different PKey values will be generated for those definitions).

Examples:

```
Default=0x7fff : ALL, SELF=full ;
Default=0x7fff : ALL, ALL_SWITCHES=full, SELF=full ;
NewPartition , ipoib : 0x123456=full, 0x3456789034=limi, 0x2134af2306 ;
YetAnotherOne = 0x300 : SELF=full ;
```

```
YetAnotherOne = 0x300 : ALL=limited ;
ShareIO = 0x80, defmember=full : 0x123451, 0x123452;
# 0x123453, 0x123454 will be limited
ShareIO = 0x80 : 0x123453, 0x123454, 0x123455=full;
# 0x123456, 0x123457 will be limited
ShareIO = 0x80 : defmember=limited : 0x123456, 0x123457, 0x123458=full;
ShareIO = 0x80, defmember=full : 0x123459, 0x12345a;
ShareIO = 0x80, defmember=full : 0x12345b, 0x12345c=limited, 0x12345d;
# multicast groups added to default
Default=0x7fff, ipoib:
mgid=ff12:401b::0707,sl=1 # random IPv4 group
mgid=ff12:601b::16 # MLDv2-capable routers
mgid=ff12:401b::16 # IGMP
mgid=ff12:601b::2 # All routers
mgid=ff12::1,sl=1,Q_Key=0xDEADBEEF,rate=3,mtu=2 # random group
ALL=full:
```

The following rule is equivalent to how OpenSM used to run prior to the partition manager:

```
Default=0x7fff,ipoib:ALL=full;
```

Effect of Topology Changes

If a link is added or removed, OpenSM may not recalculate the routes that do not have to change. A route has to change if the port is no longer UP or no longer the MinHop. When routing changes are performed, the same algorithm for balancing the routes is invoked.

In the case of using the file-based routing, any topology changes are currently ignored. The 'file' routing engine just loads the LFTs from the file specified, with no reaction to real topology. Obviously, this will not be able to recheck LIDs (by GUID) for disconnected nodes, and LFTs for non-existent switches will be skipped. Multicast is not affected by 'file' routing engine (this uses min hop tables).

Routing Algorithms

OpenSM offers the following routing engines:

1. Min Hop Algorithm

Based on the minimum hops to each node where the path length is optimized.

2. UPDN Algorithm

Based on the minimum hops to each node, but it is constrained to ranking rules. This algorithm should be chosen if the subnet is not a pure Fat Tree, and a deadlock may occur due to a loop in the subnet.

3. Fat-tree Routing Algorithm

This algorithm optimizes routing for a congestion-free "shift" communication pattern. It should be chosen if a subnet is a symmetrical Fat Tree of various types, not just a K-ary-N-Tree: non-constant K, not fully staffed, and for any CBB ratio. Similar to UPDN, Fat Tree routing is constrained to ranking rules.

4. DOR Routing Algorithm

Based on the Min Hop algorithm, but avoids port equalization except for redundant links between the same two switches. This provides deadlock free routes for hypercubes when the fabric is cabled as a hypercube and for meshes when cabled as a mesh.

5. Torus-2QoS Routing Algorithm

Based on the DOR Unicast routing algorithm specialized for 2D/3D torus topologies. Torus- 2QoS provides deadlock-free routing while supporting two quality of service (QoS) levels. Additionally, it can route around multiple failed fabric links or a single failed fabric switch without introducing deadlocks, and without changing path SL values granted before the failure.

6. Routing Chains

Allows routing configuration of different parts of a single InfiniBand subnet by different routing engines. In the current release, minhop/updn/ftree/dor/torus-2QoS/pqft can be combined.



Please note that LASH Routing Algorithm is not supported.

MINHOP/UPDN/DOR routing algorithms are comprised of two stages:

- MinHop matrix calculation. How many hops are required to get from each port to each LID. The algorithm to fill these tables is different if you run standard (min hop) or Up/Down. For standard routing, a "relaxation" algorithm is used to propagate min hop from every destination LID through neighbor switches. For Up/Down routing, a BFS from every target is used. The BFS tracks link direction (up or down) and avoid steps that will perform up after a down step was used.
- 2. Once MinHop matrices exist, each switch is visited and for each target LID a decision is made as to what port should be used to get to that LID. This step is common to standard and Up/Down routing. Each port has a counter counting the number of target LIDs going through it. When there are multiple alternative ports with same MinHop to a LID, the one with less previously assigned ports is selected.

If LMC > 0, more checks are added. Within each group of LIDs assigned to same target port:

- 1. Use only ports which have same MinHop
- 2. First prefer the ones that go to different systemImageGuid (then the previous LID of the same LMC group)
- 3. If none, prefer those which go through another NodeGuid
- 4. Fall back to the number of paths method (if all go to same node).

Min Hop Algorithm

The Min Hop algorithm is invoked by default if no routing algorithm is specified. It can also be invoked by specifying '-R minhop'.

The Min Hop algorithm is divided into two stages: computation of min-hop tables on every switch and LFT output port assignment. Link subscription is also equalized with the ability to override based on port GUID. The latter is supplied by:

```
-i <equalize-ignore-guids-file>
-ignore-guids <equalize-ignore-guids-file>
```

This option provides the means to define a set of ports (by GUIDs) that will be ignored by the link load equalization algorithm.

LMC awareness routes based on a (remote) system or on a switch basis.

UPDN Algorithm

The UPDN algorithm is designed to prevent deadlocks from occurring in loops of the subnet. A loop-deadlock is a situation in which it is no longer possible to send data between any two hosts connected through the loop. As such, the UPDN routing algorithm should be sent if the subnet is not a pure Fat Tree, and one of its loops may experience a deadlock (due, for example, to high pressure).

The UPDN algorithm is based on the following main stages:

 Auto-detect root nodes - based on the CA hop length from any switch in the subnet, a statistical histogram is built for each switch (hop num vs the number of occurrences). If the histogram reflects a specific column (higher than others) for a certain node, then it is marked as a root node. Since the algorithm is statistical, it may not find any root nodes. The list of the root nodes found by this auto-detect stage is used by the ranking process stage.

(i) Note

The user can override the node list manually.

j) Note

If this stage cannot find any root nodes, and the user did not specify a GUID list file, OpenSM defaults back to the Min Hop

routing algorithm.

- 2. Ranking process All root switch nodes (found in stage 1) are assigned a rank of 0. Using the BFS algorithm, the rest of the switch nodes in the subnet are ranked incrementally. This ranking aids in the process of enforcing rules that ensure loop-free paths.
- 3. Min Hop Table setting after ranking is done, a BFS algorithm is run from each (CA or switch) node in the subnet. During the BFS process, the FDB table of each switch node traversed by BFS is updated, in reference to the starting node, based on the ranking rules and GUID values.

At the end of the process, the updated FDB tables ensure loop-free paths through the subnet.

UPDN Algorithm Usage

Activation through OpenSM:

- Use '-R updn' option (instead of old '-u') to activate the UPDN algorithm.
- Use '-a <root_guid_file>' for adding an UPDN GUID file that contains the root nodes for ranking. If the `-a' option is not used, OpenSM uses its auto-detect root nodes algorithm.

Notes on the GUID list file:

- A valid GUID file specifies one GUID in each line. Lines with an invalid format will be discarded
- The user should specify the root switch GUIDs

Fat-tree Routing Algorithm

The fat-tree algorithm optimizes routing for "shift" communication pattern. It should be chosen if a subnet is a symmetrical or almost symmetrical fat-tree of various types. It supports not just K- ary-N-Trees, by handling for non-constant K, cases where not all leafs (CAs) are present, any Constant Bisectional Ratio (CBB)ratio. As in UPDN, fat-tree also prevents credit-loop-dead- locks.

If the root GUID file is not provided ('a' or '-root_guid_file' options), the topology has to be pure fat-tree that complies with the following rules:

- Tree rank should be between two and eight (inclusively)
- Switches of the same rank should have the same number of UP-going port groups, unless they are root switches, in which case the shouldn't have UP-going ports at all.

Note: Ports that are connected to the same remote switch are referenced as 'port group'.

- Switches of the same rank should have the same number of DOWN-going port groups, unless they are leaf switches.
- Switches of the same rank should have the same number of ports in each UP-going port group.
- Switches of the same rank should have the same number of ports in each DOWN-going port group.
- All the CAs have to be at the same tree level (rank).

If the root GUID file is provided, the topology does not have to be pure fat-tree, and it should only comply with the following rules:

- Tree rank should be between two and eight (inclusively)
- All the Compute Nodes have to be at the same tree level (rank). Note that noncompute node CAs are allowed here to be at different tree ranks.

Note: List of compute nodes (CNs) can be specified using '-u' or '--cn_guid_file' OpenSM options.

Topologies that do not comply cause a fallback to min-hop routing. Note that this can also occur on link failures which cause the topology to no longer be a "pure" fat-tree.

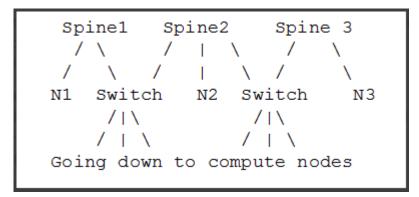
Note that although fat-tree algorithm supports trees with non-integer CBB ratio, the routing will not be as balanced as in case of integer CBB ratio. In addition to this, although the algorithm allows leaf switches to have any number of CAs, the closer the tree is to be fully populated, the more effective the "shift" communication pattern will be. In general, even if the root list is provided, the closer the topology to a pure and symmetrical fat-tree, the more optimal the routing will be.

The algorithm also dumps the compute node ordering file (opensm-ftree-ca-order.dump) in the same directory where the OpenSM log resides. This ordering file provides the CN

order that may be used to create efficient communication pattern, that will match the routing tables.

Routing between non-CN Nodes

The use of the io_guid_file option allows non-CN nodes to be located on different levels in the fat tree. In such case, it is not guaranteed that the Fat Tree algorithm will route between two non-CN nodes. In the scheme below, N1, N2, and N3 are non-CN nodes. Although all the CN have routes to and from them, there will not necessarily be a route between N1,N2 and N3. Such routes would require to use at least one of the switches the wrong way around.



To solve this problem, a list of non-CN nodes can be specified by \'-G\' or \'--io_guid_file\' option. These nodes will be allowed to use switches the wrong way around a specific number of times (specified by \'-H\' or \'--max_reverse_hops\'. With the proper max_reverse_hops and io_guid_file values, you can ensure full connectivity in the Fat Tree. In the scheme above, with a max_reverse_hop of 1, routes will be instantiated between N1<->N2 and N2<->N3. With a max_reverse_hops value of 2, N1,N2 and N3 will all have routes between them.

j) Note

Using max_reverse_hops creates routes that use the switch in a counter-stream way. This option should never be used to connect nodes with high bandwidth traffic between them! It should only be used to allow connectivity for HA purposes or similar. Also having routes the other way around can cause credit loops.

Activation through OpenSM

Use '-R ftree' option to activate the fat-tree algorithm.

i) Note

LMC > 0 is not supported by fat-tree routing. If this is specified, the default routing algorithm is invoked instead.

DOR Routing Algorithm

The Dimension Order Routing algorithm is based on the Min Hop algorithm and so uses shortest paths. Instead of spreading traffic out across different paths with the same shortest distance, it chooses among the available shortest paths based on an ordering of dimensions. Each port must be consistently cabled to represent a hypercube dimension or a mesh dimension. Paths are grown from a destination back to a source using the lowest dimension (port) of available paths at each step. This provides the ordering necessary to avoid deadlock. When there are multiple links between any two switches, they still represent only one dimension and traffic is balanced across them unless port equalization is turned off. In the case of hypercubes, the same port must be used throughout the fabric to represent the hypercube dimension and match on both ends of the cable. In the case of meshes, the dimension should consistently use the same pair of ports, one port on one end of the cable, and the other port on the other end, continuing along the mesh dimension.

Use '-R dor' option to activate the DOR algorithm.

Torus-2QoS Routing Algorithm

Torus-2QoS is a routing algorithm designed for large-scale 2D/3D torus fabrics. The torus-2QoS routing engine can provide the following functionality on a 2D/3D torus:

- Free of credit loops routing
- Two levels of QoS, assuming switches support 8 data VLs
- Ability to route around a single failed switch, and/or multiple failed links, without:

- introducing credit loops
- changing path SL values
- Very short run times, with good scaling properties as fabric size increases

Unicast Routing

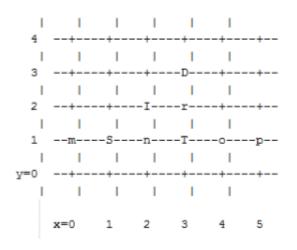
Torus-2 QoS is a DOR-based algorithm that avoids deadlocks that would otherwise occur in a torus using the concept of a dateline for each torus dimension. It encodes into a path SL which datelines the path crosses as follows:

sl = 0; for (d = 0; d < torus_dimensions; d++) /* path_crosses_dateline(d) returns 0 or 1 */ sl |= path_crosses_dateline(d) << d;</pre>

For a 3D torus, that leaves one SL bit free, which torus-2 QoS uses to implement two QoS levels. Torus-2 QoS also makes use of the output port dependence of switch SL2VL maps to encode into one VL bit the information encoded in three SL bits. It computes in which torus coordinate direc- tion each inter-switch link "points", and writes SL2VL maps for such ports as follows:

for (sl = 0; sl < 16; sl ++)
/* cdir(port) reports which torus coordinate direction a switch port
* "points" in, and returns 0, 1, or 2 */
sl2vl(iport,oport,sl) = 0x1 & (sl >> cdir(oport));

Thus, on a pristine 3D torus, i.e., in the absence of failed fabric switches, torus-2 QoS consumes 8 SL values (SL bits 0-2) and 2 VL values (VL bit 0) per QoS level to provide deadlock-free routing on a 3D torus. Torus-2 QoS routes around link failure by "taking the long way around" any 1D ring interrupted by a link failure. For example, consider the 2D 6x5 torus below, where switches are denoted by [+a-zA-Z]:



For a pristine fabric the path from S to D would be S-n-T-r-D. In the event that either link S-n or n-T has failed, torus-2QoS would use the path S-m-p-o-T-r-D.

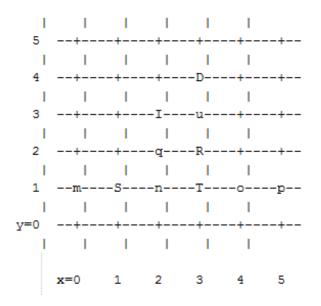
Note that it can do this without changing the path SL value; once the 1D ring m-S-n-T-op-m has been broken by failure, path segments using it cannot contribute to deadlock, and the x-direction dateline (between, say, x=5 and x=0) can be ignored for path segments on that ring. One result of this is that torus-2QoS can route around many simultaneous link failures, as long as no 1D ring is broken into disjoint segments. For example, if links n-T and T-o have both failed, that ring has been broken into two disjoint segments, T and o-pm-S-n. Torus-2QoS checks for such issues, reports if they are found, and refuses to route such fabrics.

Note that in the case where there are multiple parallel links between a pair of switches, torus-2QoS will allocate routes across such links in a round-robin fashion, based on ports at the path destination switch that are active and not used for inter-switch links. Should a link that is one of several such parallel links fail, routes are redistributed across the remaining links. When the last of such a set of parallel links fails, traffic is rerouted as described above.

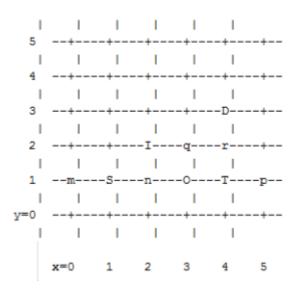
Handling a failed switch under DOR requires introducing into a path at least one turn that would be otherwise "illegal", i.e. not allowed by DOR rules. Torus-2QoS will introduce such a turn as close as possible to the failed switch in order to route around it. n the above example, suppose switch T has failed, and consider the path from S to D. Torus-2QoS will produce the path S-n-I-r-D, rather than the S-n-T-r-D path for a pristine torus, by introducing an early turn at n. Normal DOR rules will cause traffic arriving at switch I to be forwarded to switch r; for traffic arriving from I due to the "early" turn at n, this will generate an "illegal" turn at I.

Torus-2QoS will also use the input port dependence of SL2VL maps to set VL bit 1 (which would be otherwise unused) for y-x, z-x, and z-y turns, i.e., those turns that are illegal under DOR. This causes the first hop after any such turn to use a separate set of VL

values, and prevents deadlock in the presence of a single failed switch. For any given path, only the hops after a turn that is illegal under DOR can contribute to a credit loop that leads to deadlock. So in the example above with failed switch T, the location of the illegal turn at I in the path from S to D requires that any credit loop caused by that turn must encircle the failed switch at T. Thus the second and later hops after the illegal turn at I (i.e., hop r-D) cannot contribute to a credit loop because they cannot be used to construct a loop encircling T. The hop I-r uses a separate VL, so it cannot contribute to a credit loop encircling T. Extending this argument shows that in addition to being capable of routing around a single switch failure without introducing deadlock, torus-2QoS can also route around multiple failed switches on the condition they are adjacent in the last dimension routed by DOR. For example, consider the following case on a 6x6 2D torus:



Suppose switches T and R have failed, and consider the path from S to D. Torus-2QoS will generate the path S-n-q-l-u-D, with an illegal turn at switch I, and with hop l-u using a VL with bit 1 set. As a further example, consider a case that torus-2QoS cannot route without deadlock: two failed switches adjacent in a dimension that is not the last dimension routed by DOR; here the failed switches are O and T:

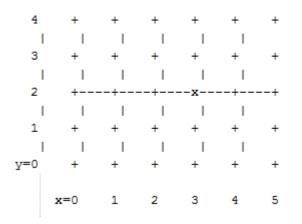


In a pristine fabric, torus-2QoS would generate the path from S to D as S-n-O-T-r-D. With failed switches O and T, torus-2QoS will generate the path S-n-I-q-r-D, with an illegal turn at switch I, and with hop I-q using a VL with bit 1 set. In contrast to the earlier examples, the second hop after the illegal turn, q-r, can be used to construct a credit loop encircling the failed switches.

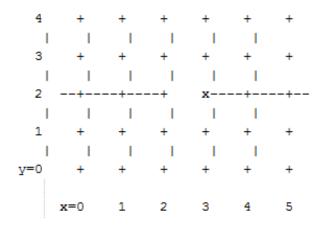
Multicast Routing

Since torus-2QoS uses all four available SL bits, and the three data VL bits that are typically available in current switches, there is no way to use SL/VL values to separate multicast traffic from unicast traffic. Thus, torus-2QoS must generate multicast routing such that credit loops cannot arise from a combination of multicast and unicast path segments. It turns out that it is possible to construct spanning trees for multicast routing that have that property. For the 2D 6x5 torus

example above, here is the full-fabric spanning tree that torus-2QoS will construct, where "x" is the root switch and each "+" is a non-root switch:

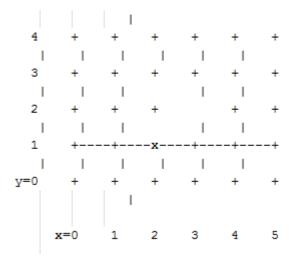


For multicast traffic routed from root to tip, every turn in the above spanning tree is a legal DOR turn. For traffic routed from tip to root, and some traffic routed through the root, turns are not legal DOR turns. However, to construct a credit loop, the union of multicast routing on this spanning tree with DOR unicast routing can only provide 3 of the 4 turns needed for the loop. In addition, if none of the above spanning tree branches crosses a dateline used for unicast credit loop avoidance on a torus, and if multicast traffic is confined to SL 0 or SL 8 (recall that torus-2QoS uses SL bit 3 to differentiate QoS level), then multicast traffic also cannot contribute to the "ring" credit loops that are otherwise possible in a torus. Torus-2QoS uses these ideas to create a master spanning tree. Every multicast group spanning tree will be constructed as a subset of the master tree, with the same root as the master tree. Such multicast group spanning trees will in general not be optimal for groups which are a subset of the full fabric. However, this compromise must be made to enable support for two QoS levels on a torus while preventing credit loops. In the presence of link or switch failures that result in a fabric for which torus-2QoS can generate credit-loop-free unicast routes, it is also possible to generate a master spanning tree for multicast that retains the required properties. For example, consider that same 2D 6x5 torus, with the link from (2,2) to (3,2) failed. Torus-2QoS will generate the following master spanning tree:



Two things are notable about this master spanning tree. First, assuming the x dateline was between x=5 and x=0, this spanning tree has a branch that crosses the dateline. However,

just as for unicast, crossing a dateline on a 1D ring (here, the ring for y=2) that is broken by a failure cannot contribute to a torus credit loop. Second, this spanning tree is no longer optimal even for multicast groups that encompass the entire fabric. That, unfortunately, is a compromise that must be made to retain the other desirable properties of torus-2QoS routing. In the event that a single switch fails, torus-2QoS will generate a master spanning tree that has no "extra" turns by appropriately selecting a root switch. In the 2D 6x5 torus example, assume now that the switch at (3,2) (i.e., the root for a pristine fabric), fails. Torus-2QoS will generate the following master spanning tree for that case:



Assuming the dateline was between y=4 and y=0, this spanning tree has a branch that crosses a dateline. However, this cannot contribute to credit loops as it occurs on a 1D ring (the ring for x=3) that is broken by failure, as in the above example.

Torus Topology Discovery

The algorithm used by torus-2QoS to construct the torus topology from the undirected graph representing the fabric requires that the radix of each dimension be configured via torus-2QoS.conf. It also requires that the torus topology be "seeded"; for a 3D torus this requires configuring four switches that define the three coordinate directions of the torus. Given this starting information, the algorithm is to examine the cube formed by the eight switch locations bounded by the corners (x,y,z) and (x+1,y+1,z+1). Based on switches already placed into the torus topology at some of these locations, the algorithm examines 4-loops of inter-switch links to find the one that is consistent with a face of the cube of switch locations and adds its switches to the discovered topology in the correct locations.

Because the algorithm is based on examining the topology of 4-loops of links, a torus with one or more radix-4 dimensions requires extra initial seed configuration. See torus-2QoS.conf(5) for details. Torus-2QoS will detect and report when it has an insufficient configuration for a torus with radix-4 dimensions.

In the event the torus is significantly degraded, i.e., there are many missing switches or links, it may happen that torus-2QoS is unable to place into the torus some switches and/or links that were discovered in the fabric, and will generate a warning in that case. A similar condition occurs if torus-2QoS is misconfigured, i.e., the radix of a torus dimension as configured does not match the radix of that torus dimension as wired, and many switches/links in the fabric will not be placed into the torus.

Quality Of Service Configuration

OpenSM will not program switches and channel adapters with SL2VL maps or VL arbitration configuration unless it is invoked with -Q. Since torus-2QoS depends on such functionality for correct operation, always invoke OpenSM with -Q when torus-2QoS is in the list of routing engines. Any quality of service configuration method supported by OpenSM will work with torus-2QoS, subject to the following limitations and considerations. For all routing engines supported by OpenSM except torus-2QoS, there is a one-to-one correspondence between QoS level and SL. Torus-2QoS can only support two quality of service levels, so only the high-order bit of any SL value used for unicast QoS configuration will be honored by torus-2QoS. For multicast QoS configuration, only SL values 0 and 8 should be used with torus-2QoS.

Since SL to VL map configuration must be under the complete control of torus-2QoS, any configuration via qos_sl2vl, qos_swe_sl2vl, etc., must and will be ignored, and a warning will be generated. Torus-2QoS uses VL values 0-3 to implement one of its supported QoS levels, and VL values 4-7 to implement the other. Hard-to-diagnose application issues may arise if traffic is not delivered fairly across each of these two VL ranges. Torus-2QoS will detect and warn if VL arbitration is configured unfairly across VLs in the range 0-3, and also in the range 4-7. Note that the default OpenSM VL arbitration configuration does not meet this constraint, so all torus-2QoS users should configure VL arbitration via qos_vlarb_high, qos_vlarb_low, etc.

Operational Considerations

Any routing algorithm for a torus IB fabric must employ path SL values to avoid credit loops. As a result, all applications run over such fabrics must perform a path record query to obtain the correct path SL for connection setup. Applications that use rdma_cm for connection setup will automatically meet this requirement.

If a change in fabric topology causes changes in path SL values required to route without credit loops, in general, all applications would need to repath to avoid message deadlock. Since torus- 2QoS has the ability to reroute after a single switch failure without changing path SL values, repathing by running applications is not required when the fabric is routed with torus-2QoS.

Torus-2QoS can provide unchanging path SL values in the presence of subnet manager failover provided that all OpenSM instances have the same idea of dateline location. See torus- 2QoS.conf(5) for details. Torus-2QoS will detect configurations of failed switches and links that prevent routing that is free of credit loops and will log warnings and refuse to route. If "no_fall- back" was configured in the list of OpenSM routing engines, then no other routing engine will attempt to route the fabric. In that case, all paths that do not transit the failed components will continue to work, and the subset of paths that are still operational will continue to remain free of credit loops. OpenSM will continue to attempt to route the fabric after every sweep interval and after any change (such as a link up) in the fabric topology. When the fabric components are repaired, full functionality will be restored. In the event OpenSM was configured to allow some other engine to route the fabric if torus-2QoS fails, then credit loops and message deadlock are likely if torus-2QoS had previously routed the fabric successfully. Even if the other engine is capable of routing a torus without credit loops, applications that built connections with path SL values granted under torus-2QoS will likely experience message deadlock under routing generated by a different engine, unless they repath. To verify that a torus fabric is routed free of credit loops, use ibdmchk to analyze data collected via ibdiagnet -vlr.

Torus-2QoS Configuration File Syntax

The file torus-2QoS.conf contains configuration information that is specific to the OpenSM routing engine torus-2QoS. Blank lines and lines where the first non-whitespace character is "#" are ignored. A token is any contiguous group of non-whitespace characters. Any tokens on a line following the recognized configuration tokens described below are ignored.

[torus|mesh] x_radix[m|M|t|T] y_radix[m|M|t|T] z_radix[m|M|t|T]

Either torus or mesh must be the first keyword in the configuration and sets the topology that torus-2QoS will try to construct. A 2D topology can be configured by specifying one of x_radix, y_radix, or z_radix as 1. An individual dimension can be configured as mesh (open) or torus (looped) by suffixing its radix specification with one of m, M, t, or T. Thus, "mesh 3T 4 5" and "torus 3 4M 5M" both specify the same topology.

Note that although torus-2QoS can route mesh fabrics, its ability to route around failed components is severely compromised on such fabrics. A failed fabric components very likely to cause a disjoint ring; see UNICAST ROUTING in torus-2QoS(8).

xp_link	sw0_GUID	sw1_GUID
yp_link	sw0_GUID	sw1_GUID
zp_link	sw0_GUID	sw1_GUID
xm_link	sw0_GUID	sw1_GUID
ym_link	sw0_GUID	sw1_GUID
zm_link	sw0_GUID	sw1_GUID

These keywords are used to seed the torus/mesh topology. For example, "xp_link 0x2000 0x2001" specifies that a link from the switch with node GUID 0x2000 to the switch with node GUID 0x2001 would point in the positive x direction, while "xm_link 0x2000 0x2001" specifies that a link from the switch with node GUID 0x2000 to the switch with node GUID 0x2001 would point in the negative x direction. All the link keywords for a given seed must specify the same "from" switch.

In general, it is not necessary to configure both the positive and negative directions for a given coordinate; either is sufficient. However, the algorithm used for topology discovery needs extra information for torus dimensions of radix four (see TOPOLOGY DISCOVERY in torus-2QoS(8)). For such cases, both the positive and negative coordinate directions must be specified.

Based on the topology specified via the torus/mesh keyword, torus-2QoS will detect and log when it has insufficient seed configuration.

GUIDx_dateline position y_dateline position z_dateline position

In order for torus-2QoS to provide the guarantee that path SL values do not change under any conditions for which it can still route the fabric, its idea of dateline position must not change relative to physical switch locations. The dateline keywords provide the means to configure such behavior.

The dateline for a torus dimension is always between the switch with coordinate 0 and the switch with coordinate radix-1 for that dimension. By default, the common switch in a torus seed is taken as the origin of the coordinate system used to describe switch location. The position parameter for a dateline keyword moves the origin (and hence the dateline) the specified amount relative to the common switch in a torus seed.

If any of the switches used to specify a seed were to fail torus-2QoS would be unable to complete topology discovery successfully. The next_seed keyword specifies that the following link and dateline keywords apply to a new seed specification.

For maximum resiliency, no seed specification should share a switch with any other seed specification. Multiple seed specifications should use dateline configuration to ensure that torus-2QoS can grant path SL values that are constant, regardless of which seed was used to initiate topology discovery.

portgroup_max_ports max_ports - This keyword specifies the maximum number of parallel inter-switch links, and also the maximum number of host ports per switch, that torus-2QoS can accommodate. The default value is 16. Torus-2QoS will log an error message during topology discovery if this parameter needs to be increased. If this keyword appears multiple times, the last instance prevails.

port_order p1 p2 p3 ... - This keyword specifies the order in which CA ports on a destination switch are visited when computing routes. When the fabric contains switches connected with multiple parallel links, routes are distributed in a round-robin fashion across such links, and so changing the order that CA ports are visited changes the distribution of routes across such links. This may be advantageous for some specific traffic patterns.

The default is to visit CA ports in increasing port order on destination switches. Duplicate values in the list will be ignored.

Example:

```
# Look for a 2D (since x radix is one) 4x5 torus.
torus 1 4 5
# y is radix-4 torus dimension, need both
# ym_link and yp_link configuration.
yp_link 0x200000 0x200005 # sw @ y=0,z=0 -> sw @ y=1,z=0
ym_link 0x200000 0x20000f # sw @ y=0,z=0 -> sw @ y=3,z=0
# z is not radix-4 torus dimension, only need one of
# zm_link or zp_link configuration.
zp_link 0x20000 0x20001 # sw @ y=0,z=0 -> sw @ y=0,z=1
```

```
next_seed
yp_link 0x20000b 0x200010 # sw @ y=2,Z=1 -> sw @ y=3,Z=1
ym_link 0x20000b 0x20000c # sw @ y=2,Z=1 -> sw @ y=1,Z=1
Zp_link 0x20000b 0x20000c # sw @ y=2,Z=1 -> sw @ y=2,Z=2
y_dateline -2 # Move the dateline for this seed
Z_dateline -1 # back to its original position.
# If OpenSM failover is configured, for maximum resiliency
# one instance should run on a host attached to a switch
# from the first seed, and another instance should run
# on a host attached to a switch from the second seed.
# Both instances should use this torus-2QoS.conf to ensure
# path SL values do not change in the event of SM failover.
# port_order defines the order on which the ports would be
# chosen for routing.
port_order 7 10 8 11 9 12 25 28 26 29 27 30
```

Routing Chains

The routing chains feature is offering a solution that enables one to configure different parts of the fabric and define a different routing engine to route each of them. The routings are done in a sequence (hence the name "chains") and any node in the fabric that is configured in more than one part is left with the routing updated by the last routing engine it was a part of.

Configuring Routing Chains

To configure routing chains:

- 1. Define the port groups.
- 2. Define topologies based on previously defined port groups.
- 3. Define configuration files for each routing engine.
- 4. Define routing engine chains over previously defined topologies and configuration files.

Defining Port Groups

The basic idea behind the port groups is the ability to divide the fabric into sub-groups and give each group an identifier that can be used to relate to all nodes in this group. The port groups is a separate feature from the routing chains but is a mandatory prerequisite for it. In addition, it is used to define the participants in each of the routing algorithms.

Defining a Port Group Policy File

In order to define a port group policy file, set the parameter 'pgrp_policy_file' in the OpenSM configuration file.

pgrp_policy_file /etc/opensm/conf/port_groups_policy_file

Configuring a Port Group Policy

The port groups policy file details the port groups in the fabric. The policy file should be composed of one or more paragraphs that define a group. Each paragraph should begin with the line 'port-group' and end with the line 'end-port-group'.

For example:

```
port-group
...port group qualifiers...
end-port-group
```

Port Group Qualifiers



Unlike the port group's beginning and end which do not require a colon, all qualifiers must end with a colon (':'). Also - a colon is a

predefined mark that must not be used inside qualifier values. The inclusion of a colon in the name or the use of a port group will result in the policy's failure.

Rule Qualifier

Parame ter	Description	Example
name	Each group must have a name. Without a name qualifier, the policy fails.	name: grp1
use	'use' is an optional qualifier that one can define in order to describe the usage of this port group (if undefined, an empty string is used as a default).	use: first port group

There are several qualifiers used to describe a rule that determines which ports will be added to the group. Each port group may include one or more rules out of the rules described in the below table (at least one rule must be defined for each port group).

Param eter	Description	Example
	Comma separated list of GUIDs to include in the group. If no specific physical ports were configured, all physical ports of the guid are chosen. However, for each guid, one can detail specific physical ports to be included in the group. This can be done using the following syntax:	
guid list	 Specify a specific port in a guid to be chosen port-guid: 0x283@3 Specify a specific list of ports in a guid to be chosen port-guid: 0x286@1/5/7 	port-guid: 0x283,
	 Specify a specific range of ports in a guid to be chosen port-guid: 0x289@2-5 	0x286, 0x289
	 Specify a list of specific ports and ports ranges in a guid to be chosen port-guid: 0x289@2-5/7/9-13/18 	
	 Complex rule port-guid: 0x283@5-8/12/14, 0x286, 0x289/6/ 8/12 	

Param eter	Description	Example
port guid range	It is possible to configure a range of guids to be chosen to the group. However, while using the range qualifier, it is impossible to detail specific physical ports. Note: A list of ranges cannot be specified. The below example is invalid and will cause the policy to fail: port-guid-range: 0x283-0x289, 0x290- 0x295	port-guid-range: 0x283-0x289
port name	One can configure a list of hostnames as a rule. Hosts with a node description that is built out of these hostnames will be chosen. Since the node description contains the network card index as well, one might also specify a network card index and a physical port to be chosen. For example, the given configuration will cause only physical port 2 of a host with the node description 'kuku HCA-1' to be chosen. port and hca_idx parameters are optional. If the port is unspecified, all physical ports are chosen. If hca_idx is unspecified, all card numbers are chosen. Specifying a hostname is mandatory. One can configure a list of hostname/ port/hca_idx sets in the same qualifier as follows: port-name: hostname=kuku; port=2; hca_idx=1, hostname=host1; port=3, hostname=host2 Note: port-name qualifier is not relevant for switches, but for HCA's only.	port-name: host- name=kuku; port=2; hca_idx=1
port regexp	One can define a regular expression so that only nodes with a matching node description will be chosen to the group. Note: This example shows how to choose nodes which their node description starts with 'SW'.	port-regexp: SW
	It is possible to specify one physical port to be chosen for matching nodes (there is no option to define a list or a range of ports). The given example will cause only nodes that match physical port 3 to be added to the group.	port-regexp: SW:3
union rule	It is possible to define a rule that unites two different port groups. This means that all ports from both groups will be included in the united group.	union-rule: grp1, grp2
subtract rule	One can define a rule that subtracts one port group from another. The given rule, for example, will cause all the ports which are a part of grp1, but not included in grp2, to be chosen. In subtraction (unlike union), the order does matter, since the purpose is to subtract the second group from the first one.	subtract-rule: grp1, grp2

Param eter	Description	Example
	There is no option to define more than two groups for union/subtraction. However, one can unite/subtract groups which are a union or a subtraction themselves, as shown in the port groups policy file example.	

Predefined Port Groups

There are 3 predefined, automatically created port groups that are available for use, yet cannot be defined in the policy file (if a group in the policy is configured with the name of one of these predefined groups, the policy fails) -

- ALL a group that includes all nodes in the fabric
- ALL_SWITCHES a group that includes all switches in the fabric
- ALL_CAS a group that includes all HCAs in the fabric
- ALL_ROUTERS a group that includes all routers in the fabric (supported in OpenSM starting from v4.9.0)

Port Groups Policy Examples

```
port-group
name: grp3
use: Subtract of groups grp1 and grp2
subtract-rule: grp1, grp2
end-port-group
port-group
name: grp1
port-guid: 0x281, 0x282, 0x283
end-port-group
```

```
port-group
name: grp2
port-guid-range: 0x282-0x286
port-name: hostname=server1 port=1
end-port-group
name: grp4
port-name: hostname=kika port=1 hca_idx=1
end-port-group
name: grp3
union-rule: grp3, grp4
end-port-group
```

Defining a Topologies Policy File

In order to define a topology policy file, set the parameter 'topo_policy_file' in the OpenSM configuration file.

topo_policy_file /etc/opensm/conf/topo_policy_file.cfg

Configuring a Topology Policy

The topologies policy file details a list of topologies. The policy file should be composed of one or more paragraphs which define a topology. Each paragraph should begin with the line 'topol- ogy' and end with the line 'end-topology'.

For example:

```
topology
...topology qualifiers...
```

Topology Qualifiers

i) Note

Unlike topology and end-topology which do not require a colon, all qualifiers must end with a colon (':'). Also - a colon is a predefined mark that must not be used inside qualifier values. An inclusion of a column in the qualifier values will result in the policy's failure.

All topology qualifiers are mandatory. Absence of any of the below qualifiers will cause the policy parsing to fail.

Topology Qualifiers

Parameter	Description	Example	
id	Topology ID. Legal Values – any positive value. Must be unique.	es all switches and sw-grn: vs_switches	
sw-grp	Name of the port group that includes all switches and switch ports to be used in this topology.		
hca-grp	Name of the port group that includes all HCA's to be used in this topology.		

Configuration File per Routing Engine

Each engine in the routing chain can be provided by its own configuration file. Routing engine configuration file is the fraction of parameters defined in the main OpenSM configuration file.

Some rules should be applied when defining a particular configuration file for a routing engine:

- Parameters that are not specified in specific routing engine configuration file are inherited from the main OpenSM configuration file.
- The following configuration parameters are taking effect only in the main OpenSM configuration file:
 - qos and qos_* settings like (vl_arb, sl2vl, etc.)
 - Imc
 - routing_engine

Defining a Routing Chain Policy File

In order to define a port group policy file, set the parameter 'rch_policy_file' in the OpenSM configuration file.

rch_policy_file /etc/opensm/conf/chains_policy_file

First Routing Engine in the Chain

The first unicast engine in a routing chain must include all switches and HCAs in the fabric (topology id must be 0). The path-bit parameter value is path-bit 0 and it cannot be changed.

Configuring a Routing Chains Policy

The routing chains policy file details the routing engines (and their fallback engines) used for the fabric's routing. The policy file should be composed of one or more paragraphs which defines an

engine (or a fallback engine). Each paragraph should begin with the line 'unicast-step' and end with the line 'end-unicast-step'.

For example:

```
unicast-step
…routing engine qualifiers…
end-unicast-step
```

Routing Engine Qualifiers

(i) Note

Unlike unicast-step and end-unicast-step which do not require a colon, all qualifiers must end with a colon (':'). Also - a colon is a predefined mark that must not be used inside qualifier values. An inclusion of a colon in the qualifier values will result in the policy's failure.

Parameter	Description	Example
	'id' is mandatory. Without an ID qualifier for each engine, the policy fails.	
id	 Legal values – size_t value (0 is illegal). The engines in the policy chain are set according to an ascending id order, so it is highly crucial to verify that the id that is given to the engines match the order in which you would like the engines to be set. 	is: 1
engine	This is a mandatory qualifier that describes the routing algorithm used within this unicast step. Currently, on the first phase of routing chains, legal values are minhop/ftree/updn.	engine: minhop
use	This is an optional qualifier that enables one to describe the usage of this unicast step. If undefined, an empty string is used as a default.	use: ftree routing for for yellow stone nodes
config	This is an optional qualifier that enables one to define a separate OpenSM config file for a specific unicast step. If	config: /etc/config/

Parameter	Description	Example
	undefined, all parameters are taken from main OpenSM configuration file.	opensm2.cfg
topology	 Define the topology that this engine uses. Legal value - id of an existing topology that is defined in topologies policy (or zero that represents the entire fabric and not a specific topology). Default value - If unspecified, a routing engine will relate to the entire fabric (as if topology zero was defined). Notice: The first routing engine (the engine with the lowest id) MUST be configured with topology: 0 (entire fabric) or else, the routing chain parser will fail. 	topology: 1
fallback-to	 This is an optional qualifier that enables one to define the current unicast step as a fallback to another unicast step. This can be done by defining the id of the unicast step that this step is a fallback to. If undefined, the current unicast step is not a fallback. If the value of this qualifier is a non-existent engine id, this step will be ignored. A fallback step is meaningless if the step it is a fallback to did not fail. It is impossible to define a fallback to a fall- back step (such definition will be ignored) 	-
path-bit	This is an optional qualifier that enables one to define a specific lid offset to be used by the current unicast step. Setting Imc > 0 in main OpenSM configuration file is a prerequisite for assigning specific path-bit for the routing engine. Default value is 0 (if path-bit is not specified)	Path-bit: 1

Dump Files per Routing Engine

Each routing engine on the chain will dump its own data files if the appropriate log_flags is set (for instance 0x43).

The files that are dumped by each engine are:

• opensm-lid-matrix.dump

- opensm-lfts.dump
- opensm.fdbs
- opensm-subnet.lst

These files should contain the relevant data for each engine topology.



sl2vl and mcfdbs files are dumped only once for the entire fabric and NOT by every routing engine.

- Each engine concatenates its ID and routing algorithm name in its dump files names, as follows:
 - opensm-lid-matrix.2.minhop.dump
 - opensm.fdbs.3.ftree
 - opensm-subnet.4.updn.lst
- In case that a fallback routing engine is used, both the routing engine that failed and the fallback engine that replaces it, dump their data.

If, for example, engine 2 runs ftree and it has a fallback engine with 3 as its id that runs minhop, one should expect to find 2 sets of dump files, one for each engine:

- opensm-lid-matrix.2.ftree.dump
- opensm-lid-matrix.3.minhop.dump
- opensm.fdbs.2.ftree
- opensm.fdbs.3.munhop

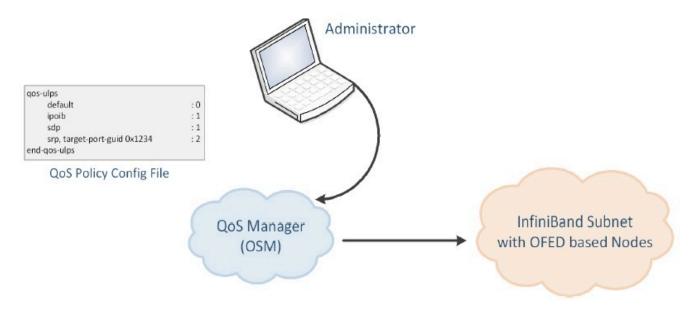
Unicast Routing Cache

Unicast routing cache prevents routing recalculation (which is a heavy task in a large cluster) when no topology change was detected during the heavy sweep, or when the topology change does not require new routing calculation (for example, when one or more CAs/RTRs/leaf switches going down, or one or more of these nodes coming back after being down).

Quality of Service Management in OpenSM

When Quality of Service (QoS) in OpenSM is enabled (using the '-Q' or '--qos' flags), OpenSM looks for a QoS Policy file. During fabric initialization and at every heavy sweep, OpenSM parses the QoS policy file, applies its settings to the discovered fabric elements, and enforces the provided policy on client requests. The overall flow for such requests is as follows:

- The request is matched against the defined matching rules such that the QoS Level definition is found
- Given the QoS Level, a path(s) search is performed with the given restrictions imposed by that level



There are two ways to define QoS policy:

- Advanced the advanced policy file syntax provides the administrator various ways to match a PathRecord/MultiPathRecord (PR/MPR) request, and to enforce various QoS constraints on the requested PR/MPR
- Simple the simple policy file syntax enables the administrator to match PR/MPR requests by various ULPs and applications running on top of these ULPs

Advanced QoS Policy File

The QoS policy file has the following sections:

- 1. Port Groups (denoted by port-groups) this section defines zero or more port groups that can be referred later by matching rules (see below). Port group lists ports by:
 - Port GUID
 - Port name, which is a combination of NodeDescription and IB port number

- PKey, which means that all the ports in the subnet that belong to partition with a given PKey belong to this port group

- Partition name, which means that all the ports in the subnet that belong to partition with a given name belong to this port group

- Node type, where possible node types are: CA, SWITCH, ROUTER, ALL, and SELF (SM's port).

- 2. QoS Setup (denoted by qos-setup) this section describes how to set up SL2VL and VL Arbitration tables on various nodes in the fabric. However, this is not supported in OFED. SL2VL and VLArb tables should be configured in the OpenSM options file (default location /var/cache/opensm/opensm.opts).
- 3. QoS Levels (denoted by qos-levels) each QoS Level defines Service Level (SL) and a few optional fields:
 - MTU limit
 - Rate limit
 - PKey
 - Packet lifetime

When path(s) search is performed, it is done with regards to restriction that these QoS Level parameters impose. One QoS level that is mandatory to define is a DEFAULT QoS level. It is applied to a PR/MPR query that does not match any existing match rule. Similar to any other QoS Level, it can also be explicitly referred by any match rule.

• QoS Matching Rules (denoted by qos-match-rules) - each PathRecord/MultiPathRecord query that OpenSM receives is matched against the set of matching rules. Rules are scanned in order of appearance in the QoS policy file such as the first match takes precedence.

Each rule has a name of QoS level that will be applied to the matching query. A default QoS level is applied to a query that did not match any rule.

Queries can be matched by:

- Source port group (whether a source port is a member of a specified group)
- Destination port group (same as above, only for destination port)
- PKey
- QoS class
- Service ID

To match a certain matching rule, PR/MPR query has to match ALL the rule's criteria. However, not all the fields of the PR/MPR query have to appear in the matching rule.

For instance, if the rule has a single criterion - Service ID, it will match any query that has this Service ID, disregarding rest of the query fields. However, if a certain query has only Service ID (which means that this is the only bit in the PR/MPR component mask that is on), it will not match any rule that has other matching criteria besides Service ID.

Simple QoS Policy Definition

Simple QoS policy definition comprises of a single section denoted by qos-ulps. Similar to the advanced QoS policy, it has a list of match rules and their QoS Level, but in this case a match rule has only one criterion - its goal is to match a certain ULP (or a certain application on top of this ULP) PR/MPR request, and QoS Level has only one constraint - Service Level (SL).

The simple policy section may appear in the policy file in combine with the advanced policy, or as a stand-alone policy definition. See more details and list of match rule criteria below.

Policy File Syntax Guidelines

- Leading and trailing blanks, as well as empty lines, are ignored, so the indentation in the example is just for better readability.
- Comments are started with the pound sign (#) and terminated by EOL.
- Any keyword should be the first non-blank in the line, unless it's a comment.
- Keywords that denote section/subsection start have matching closing keywords.
- Having a QoS Level named "DEFAULT" is a must it is applied to PR/MPR requests that did not match any of the matching rules.
- Any section/subsection of the policy file is optional.

Examples of Advanced Policy Files

As mentioned earlier, any section of the policy file is optional, and the only mandatory part of the policy file is a default QoS Level.

Here is an example of the shortest policy file:

```
qos-levels
    qos-level
    name: DEFAULT
    sl: 0
    end-qos-level
end-qos-levels
```

Port groups section is missing because there are no match rules, which means that port groups are not referred anywhere, and there is no need defining them. And since this policy file doesn't have any matching rules, PR/MPR query will not match any rule, and OpenSM will enforce default QoS level. Essentially, the above example is equivalent to not having a QoS policy file at all.

The following example shows all the possible options and keywords in the policy file and their syntax:

```
#
    # See the comments in the following example.
    # They explain different keywords and their meaning.
    #
    port-groups
        port-group # using port GUIDs
            name: Storage
            # "use" is just a description that is used for logging
            # Other than that, it is just a comment
            use: SRP Targets
            port-quid: 0x100000000001, 0x10000000005-0x100000000FFFA
            port-guid: 0x100000000FFFF
        end-port-group
        port-group
            name: Virtual Servers
            # The syntax of the port name is as follows:
                 "node_description/Pnum".
            #
            # node_description is compared to the NodeDescription
of the node,
            # and "Pnum" is a port number on that node.
            port-name: "vs1 HCA-1/P1, vs2 HCA-1/P1"
        end-port-group
        # using partitions defined in the partition policy
        port-group
            name: Partitions
            partition: Part1
            pkey: 0x1234
        end-port-group
        # using node types: CA, ROUTER, SWITCH, SELF (for node
that runs SM)
        # or ALL (for all the nodes in the subnet)
```

```
port-group
            name: CAs and SM
            node-type: CA, SELF
        end-port-group
    end-port-groups
    qos-setup
        # This section of the policy file describes how to set up
SL2VL and VL
        # Arbitration tables on various nodes in the fabric.
        # However, this is not supported in OFED - the section is
parsed
        # and ignored. SL2VL and VLArb tables should be
configured in the
        # OpenSM options file (by default -
/var/cache/opensm/opensm.opts).
    end-gos-setup
    qos-levels
        # Having a QoS Level named "DEFAULT" is a must - it is
applied to
        # PR/MPR requests that didn't match any of the matching
rules.
        qos-level
            name: DEFAULT
            use: default QoS Level
            sl: 0
        end-qos-level
        # the whole set: SL, MTU-Limit, Rate-Limit, PKey, Packet
Lifetime
        gos-level
            name: WholeSet
            sl: 1
```

```
mtu-limit: 4
            rate-limit: 5
            pkey: 0x1234
            packet-life: 8
        end-qos-level
    end-qos-levels
    # Match rules are scanned in order of their appearance in the
policy file.
    # First matched rule takes precedence.
    qos-match-rules
        # matching by single criteria: QoS class
        gos-match-rule
            use: by QoS class
            qos-class: 7-9,11
            # Name of qos-level to apply to the matching PR/MPR
            qos-level-name: WholeSet
        end-gos-match-rule
        # show matching by destination group and service id
        gos-match-rule
            use: Storage targets
            destination: Storage
            service-id: 0x100000000001, 0x10000000008-0x100000000FFF
            gos-level-name: WholeSet
        end-qos-match-rule
        gos-match-rule
            source: Storage
            use: match by source group only
            qos-level-name: DEFAULT
        end-gos-match-rule
        gos-match-rule
            use: match by all parameters
```

Simple QoS Policy - Details and Examples

Simple QoS policy match rules are tailored for matching ULPs (or some application on top of a ULP) PR/MPR requests. This section has a list of per-ULP (or per-application) match rules and the SL that should be enforced on the matched PR/MPR query.

Match rules include:

- Default match rule that is applied to PR/MPR query that didn't match any of the other match rules
- IPoIB with a default PKey
- IPoIB with a specific PKey
- Any ULP/application with a specific Service ID in the PR/MPR query
- Any ULP/application with a specific PKey in the PR/MPR query
- Any ULP/application with a specific target IB port GUID in the PR/MPR query

Since any section of the policy file is optional, as long as basic rules of the file are kept (such as no referring to nonexistent port group, having default QoS Level, etc), the simple policy section (qos-ulps) can serve as a complete QoS policy file.

The shortest policy file in this case would be as follows:

qos-ulps
 default : 0 #default SL

It is equivalent to the previous example of the shortest policy file, and it is also equivalent to not having policy file at all. Below is an example of simple QoS policy with all the possible keywords:

```
qos-ulps
                   :0 # default SL
default
sdp, port-num 30000 :0 # SL for application running on
                                   # top of SDP when a destination
                                   # TCP/IPport is 30000
sdp, port-num 10000-20000 : 0
sdp
                       :1 # default SL for any other
                                   # application running on top of
SDP
                      :2 # SL for RDS traffic
rds
ipoib, pkey 0x0001 :0 # SL for IPoIB on partition with
                                   # pkey 0x0001
                       :4 # default IPoIB partition,
ipoib
                                   # pkey=0x7FFF
any, service-id 0x6234:6 # match any PR/MPR query with a
                                   # specific Service ID
any, pkey 0x0ABC
                     :6 # match any PR/MPR guery with a
                                   # specific PKey
srp, target-port-guid 0x1234 : 5 # SRP when SRP Target is located
                                   # on a specified IB port GUID
any, target-port-guid 0x0ABC-0xFFFFF : 6 # match any PR/MPR query
                                   # with a specific target port
GUID
end-qos-ulps
```

Similar to the advanced policy definition, matching of PR/MPR queries is done in order of appearance in the QoS policy file such as the first match takes precedence, except for the "default" rule, which is applied only if the query didn't match any other rule. All other

sections of the QoS policy file take precedence over the qos-ulps section. That is, if a policy file has both qos-match-rules and qos-ulps sections, then any query is matched first against the rules in the qos-match-rules section, and only if there was no match, the query is matched against the rules in qos-ulps section.

Note that some of these match rules may overlap, so in order to use the simple QoS definition effectively, it is important to understand how each of the ULPs is matched.

IPoIB

IPoIB query is matched by PKey or by destination GID, in which case this is the GID of the multicast group that OpenSM creates for each IPoIB partition.

Default PKey for IPoIB partition is 0x7fff, so the following three match rules are equivalent:

```
ipoib:<SL>ipoib, pkey 0x7fff : <SL>
any, pkey 0x7fff : <SL>
```

SRP

Service ID for SRP varies from storage vendor to vendor, thus SRP query is matched by the target IB port GUID. The following two match rules are equivalent:

```
srp, target-port-guid 0x1234 : <SL>
any, target-port-guid 0x1234 : <SL>
```

Note that any of the above ULPs might contain target port GUID in the PR query, so in order for these queries not to be recognized by the QoS manager as SRP, the SRP match rule (or any match rule that refers to the target port GUID only) should be placed at the end of the qos-ulps match rules.

MPI

SL for MPI is manually configured by an MPI admin. OpenSM is not forcing any SL on the MPI traffic, which explains why it is the only ULP that did not appear in the qos-ulps section.

SL2VL Mapping and VL Arbitration

OpenSM cached options file has a set of QoS related configuration parameters, that are used to configure SL2VL mapping and VL arbitration on IB ports. These parameters are:

- Max VLs: the maximum number of VLs that will be on the subnet
- High limit: the limit of High Priority component of VL Arbitration table (IBA 7.6.9)
- VLArb low table: Low priority VL Arbitration table (IBA 7.6.9) template
- VLArb high table: High priority VL Arbitration table (IBA 7.6.9) template
- SL2VL: SL2VL Mapping table (IBA 7.6.6) template. It is a list of VLs corresponding to SLs 0-15 (Note that VL15 used here means drop this SL).

There are separate QoS configuration parameters sets for various target types: CAs, routers, switch external ports, and switch's enhanced port 0. The names of such parameters are prefixed by "qos_<type>_" string. Here is a full list of the currently supported sets:

- qos_ca_—QoS configuration parameters set for CAs.
- qos_rtr_—parameters set for routers.
- qos_sw0_ —parameters set for switches' port 0.
- qos_swe_ —parameters set for switches' external ports.

Here's the example of typical default values for CAs and switches' external ports (hard-coded in OpenSM initialization):

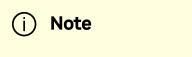
```
qos_ca_max_vls 15
qos_ca_high_limit 0
```

```
qos_ca_vlarb_high
0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0
qos_ca_vlarb_low
0:0,1:4,2:4,3:4,4:4,5:4,6:4,7:4,8:4,9:4,10:4,11:4,12:4,13:4,14:4
qos_ca_sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7
qos_swe_max_vls 15
qos_swe_high_limit 0
qos_swe_vlarb_high
0:4,1:0,2:0,3:0,4:0,5:0,6:0,7:0,8:0,9:0,10:0,11:0,12:0,13:0,14:0
qos_swe_vlarb_low
0:0,1:4,2:4,3:4,4:4,5:4,6:4,7:4,8:4,9:4,10:4,11:4,12:4,13:4,14:4
qos_swe_sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7
```

VL arbitration tables (both high and low) are lists of VL/Weight pairs. Each list entry contains a VL number (values from 0-14), and a weighting value (values 0-255), indicating the number of 64 byte units (credits) which may be transmitted from that VL when its turn in the arbitration occurs. A weight of 0 indicates that this entry should be skipped. If a list entry is programmed for VL15 or for a VL that is not supported or is not currently configured by the port, the port may either skip that entry or send from any supported VL for that entry.

Note, that the same VLs may be listed multiple times in the High or Low priority arbitration tables, and, further, it can be listed in both tables. The limit of high-priority VLArb table (qos_high_limit) indicates the number of high-priority packets that can be transmitted without an opportunity to send a low-priority packet. Specifically, the number of bytes that can be sent is high_limit times 4K bytes.

A high_limit value of 255 indicates that the byte limit is unbounded.



If the 255 value is used, the low priority VLs may be starved.

A value of 0 indicates that only a single packet from the high-priority table may be sent before an opportunity is given to the low-priority table. Keep in mind that ports usually transmit packets of size equal to MTU. For instance, for 4KB MTU a single packet will require 64 credits, so in order to achieve effective VL arbitration for packets of 4KB MTU, the weighting values for each VL should be multiples of 64.

Below is an example of SL2VL and VL Arbitration configuration on subnet:

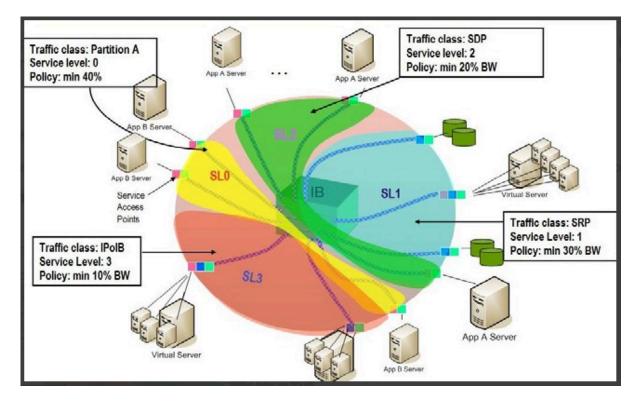
qos_ca_max_vls 15 qos_ca_high_limit 6 qos_ca_vlarb_high 0:4 qos_ca_vlarb_low 0:0,1:64,2:128,3:192,4:0,5:64,6:64,7:64 qos_ca_sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7 qos_swe_max_vls 15 qos_swe_high_limit 6 qos_swe_vlarb_high 0:4 qos_swe_vlarb_low 0:0,1:64,2:128,3:192,4:0,5:64,6:64,7:64 qos_swe_sl2vl 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,7

In this example, there are 8 VLs configured on subnet: VL0 to VL7. VL0 is defined as a high priority VL, and it is limited to 6 x 4KB = 24KB in a single transmission burst. Such configuration would suilt VL that needs low latency and uses small MTU when transmitting packets. Rest of VLs are defined as low priority VLs with different weights, while VL4 is effectively turned off.

Deployment Example

The figure below shows an example of an InfiniBand subnet that has been configured by a QoS manager to provide different service levels for various ULPs.

QoS Deployment on InfiniBand Subnet Example



QoS Configuration Examples

The following are examples of QoS configuration for different cluster deployments. Each example provides the QoS level assignment and their administration via OpenSM configuration files.

Typical HPC Example: MPI and Lustre

Assignment of QoS Levels

- MPI
 - Separate from I/O load
 - Min BW of 70%
- Storage Control (Lustre MDS)
 - Low latency
- Storage Data (Lustre OST)
 - Min BW 30%

Administration

• MPI is assigned an SL via the command line

host1# mpirun –sl 0

• OpenSM QoS policy file

Note: In this policy file example, replace OST* and MDS* with the real port GUIDs.

• OpenSM options file

```
qos_max_vls 8
qos_high_limit 0
qos_vlarb_high 2:1
qos_vlarb_low 0:96,1:224
qos_sl2vl 0,1,2,3,4,5,6,7,15,15,15,15,15,15,15,15,15
```

EDC SOA (2-tier): IPoIB and SRP

The following is an example of QoS configuration for a typical enterprise data center (EDC) with service oriented architecture (SOA), with IPoIB carrying all application traffic

and SRP used for storage.

QoS Levels

- Application traffic
 - IPoIB (UD and CM) and SDP
 - Isolated from storage
 - Min BW of 50%
- SRP
 - Min BW 50%
 - Bottleneck at storage nodes

Administration

• OpenSM QoS policy file

qos-u]	Lps	
	default	
:0		
	ipoib	
:1		
	sdp	
:1		
	<pre>srp, target-port-guid SRPT1,SRPT2,SRPT3</pre>	:2
end-qo	os-ulps	

Note: In this policy file example, replace SRPT* with the real SRP Target port GUIDs.

• OpenSM options file

```
qos_max_vls 8
qos_high_limit 0
qos_vlarb_high 1:32,2:32
qos_vlarb_low 0:1,
qos_sl2vl 0,1,2,3,4,5,6,7,15,15,15,15,15,15,15,15,15
```

EDC (3-tier): IPoIB, RDS, SRP

The following is an example of QoS configuration for an enterprise data center (EDC), with IPoIB carrying all application traffic, RDS for database traffic, and SRP used for storage.

QoS Levels

- Management traffic (ssh)
 - IPoIB management VLAN (partition A)
 - Min BW 10%
- Application traffic
 - IPoIB application VLAN (partition B)
 - Isolated from storage and database
 - Min BW of 30%
- Database Cluster traffic
 - RDS
 - Min BW of 30%
- SRP
 - Min BW 30%
 - Bottleneck at storage nodes

Administration

• OpenSM QoS policy file

```
qos-ulps
        default
:0
        ipoib, pkey 0x8001
:1
        ipoib, pkey 0x8002
:2
        rds
:3
        srp, target-port-guid SRPT1, SRPT2, SRPT3 :4
end-qos-ulps
```

Note: In the following policy file example, replace SRPT* with the real SRP Initiator port GUIDs.

• OpenSM options file

```
qos_max_vls 8
qos_high_limit 0
qos_vlarb_high 1:32,2:96,3:96,4:96
qos_vlarb_low 0:1
qos_sl2vl 0,1,2,3,4,5,6,7,15,15,15,15,15,15,15,15,15
```

• Partition configuration file

```
Default=0x7fff,ipoib : ALL=full;PartA=0x8001, sl=1, ipoib :
```

Enhanced QoS

Enhanced QoS provides a higher resolution of QoS at the service level (SL). Users can configure rate limit values per SL for physical ports, virtual ports, and port groups, using enhanced_qos_policy_file configuration parameter.

Valid values of this parameter:

- Full path to the policy file through which Enhanced QoS Manager is configured
- "null" to disable the Enhanced QoS Manager (default value)

(i) Note

To enable Enhanced QoS Manager, QoS must be enabled in OpenSM.

Enhanced QoS Policy File

The policy file is comprised of three sections:

• BW_NAMES: Used to define bandwidth setting and name (currently, rate limit is the only setting). Bandwidth names can be used in BW_RULES and VPORT_BW_RULES sections.

Bandwidth names are defined using the syntax:

<name> = <rate limit in 1Mbps units>

Example: My_bandwidth = 50

• BW_RULES: Used to define the rules that map the bandwidth setting to a specific SL of a specific GUID.

Bandwidth rules are defined using the syntax:

<guid>|<port group name> = <sl id>:<bandwidth name>, <sl id>:<bandwidth name>...

Examples:

0x2c9000000025 = 5:My_bandwidth, 7:My_bandwidth

Port_grp1 = 3:My_bandwidth, 9:My_bandwidth

• VPORT_BW_RULES: Used to define the rules that map the bandwidth setting to a specific SL of a specific virtual port GUID.

Bandwidth rules are defined using the syntax:

<guid>= <sl id>:<bandwidth name>, <sl id>:<bandwidth name>...Examples:

0x2c9000000026= 5:My_bandwidth, 7:My_bandwidth

Special Keywords

• Keyword "all" allows setting a rate limit of all SLs to some BW for a specific physical or virtual port. It is possible to combine "all" with specific SL rate limits.

Example:

0x2c9000000025 = all:BW1,SL3:BW2In this case, SL3 will be assigned BW2 rate limit, while the rest of SLs get BW1 rate limit.

• "default" is a well-known name which can be used to define a default rule used for any GUID with no defined rule.

If no default rule is defined, any GUID without a specific rule will be configured with unlimited rate limit for all SLs.

Keyword "all" is also applicable to the default rule. Default rule is local to each section.

Special Subnet Manager Configuration Options

New SM configuration option enhanced_qos_vport0_unlimit_default_rl was added to opensm.conf.

The possible values for this configuration option are:

• TRUE: For specific virtual portO GUID, SLs not mentioned in bandwidth rule will be set to unlimited bandwidth (0) regardless of the default rule of the VPORT_BW_RULES section.

Virtual portO GUIDs not mentioned in VPORT_BW_SECTION will be set to unlimited BW on all SLs.

• FALSE: The GUID of virtual portO is treated as any other virtual port in VPORT_BW_SECTION.

SM should be signaled by HUP once the option is changed.

Default: TRUE

Notes

- When rate limit is set to 0, it means that the bandwidth is unlimited.
- Any unspecified SL in a rule will be set to 0 (unlimited) rate limit automatically if no default rule is specified.
- Failure to complete policy file parsing leads to an undefined behavior. User must confirm no relevant error messages in SM log in order to ensure Enhanced QoS Manager is configured properly.
- A file with only 'BW_NAMES' and 'BW_RULES' keywords configures the network with an unlimited rate limit.
- HCA physical port GUID can be specified in BW_RULES and VPORT_BW_RULES sections.
- In BW_RULES section, the rate limit assigned to a specific SL will limit the total BW that can be sent through the PF on a given SL.
- In VPORT_BW_RULES section, the rate limit assigned to a specific SL will limit only the traffic sent from the IB interface corresponding to the physical port GUID (virtual portO IB interface). The traffic sent from other virtual IB interfaces will not be limited if no specific rules are defined.

Policy File Example

All physical ports in the fabric are with a rate limit of 50Mbps on SL1, except for GUID 0x2c9000000025, which is configured with rate limit of 25Mbps on SL1. In this example, the traffic on SLs (other than SL1) is unlimited.

All virtual ports in the fabric (except virtual portO of all physical ports) will be rate-limited to 15Mbps for all SLs because of the default rule of VPORT_BW_RULES section.

Virtual port GUID 0x2c9000000026 is configured with a rate limit of 10Mbps on SL3. The rest of the SLs on this virtual port will get a rate limit of 15 Mbps because of the default rule of VPORT_BW_RULES section.

```
BW_NAMES
bw1 = 50
bw2 = 25
bw3 = 15
bw4 = 10
BW_RULES
default= 1:bw1
0x2c900000025= 1:bw2
VPORT_BW_RULES
default= all:bw3
0x2c900000026= 3:bw4
-------
```

Adaptive Routing Manager and Self-Healing Networking

Adaptive Routing Manager supports advanced InfiniBand features; Adaptive Routing (AR) and Self-Healing Networking.

For information on how to set up AR and Self-Healing Networking, please refer to <u>HowTo</u> <u>Configure Adaptive Routing and Self-Healing Networking</u> Community post.

DOS MAD Prevention

DOS MAD prevention is achieved by assigning a threshold for each agent's RX. Agent's RX threshold provides a protection mechanism to the host memory by limiting the agents' RX with a threshold. Incoming MADs above the threshold are dropped and are not queued to the agent's RX.

To enable DOS MAD Prevention:

- 1. Go to /etc/modprobe.d/mlnx.conf.
- 2. Add to the file the option below.

ib_umad enable_rx_threshold 1

The threshold value can be controlled from the user-space via libibumad.

To change the value, use the following API:

int umad_update_threshold(int fd, int threshold);

@fd: file descriptor, agent's RX associated to this fd. @threshold: new threshold value

IB Router Support in OpenSM

In order to enable the IB router in OpenSM, the following parameters should be configured:

IB Router Parameters for OpenSM

Parameter	Description	Default Value
rtr_pr_flow_labe	Defines whether the SM should create alias GUIDs required for router support for each port. Defines flow label value to use in response for path records related to the router.	0 (Disabled)
rtr_pr_tclas s	Defines TClass value to use in response for path records related to the router	0
rtr_pr_sl	Defines sl value to use in response for path records related to router.	0
rtr_p_mtu	Defines MTU value to use in response for path records related to the router.	4 (IB_MTU_LEN_2048)
rtr_pr_rate Defines rate value to use in response for path records related to the router.		16 (IB_PATH_RE- CORD_RATE_100_GB S)

OpenSM Activity Report

OpenSM can produce an activity report in a form of a dump file which details the different activities done in the SM. Activities are divided into subjects. The OpenSM Supported Activities table below specifies the different activities currently supported in the SM activity report.

Reporting of each subject can be enabled individually using the configuration parameter activity_report_subjects:

• Valid values:

Comma separated list of subjects to dump. The current supported subjects are:

- "mc" activity IDs 1, 2 and 8
- "prtn" activity IDs 3, 4, and 5
- "virt" activity IDs 6 and 7
- "routing" activity IDs 8-12

Two predefined values can be configured as well:

•

- "all" dump all subjects
- "none" disable the feature by dumping none of the subjects
- Default value: "none"

OpenSM Supported Activities

ACtivity ID	Activity Name	Additional Fields	Comments	Description
1	mcm_member	MLidMGidPort GuidJoin State	Join state: 1 - Join -1 - Leave	Member joined/ left MC group
2	mcg_change	MLidMGidChange	Change: 0 - Create 1 - Delete	MC group created/deleted
3	prtn_guid_add	Port GuidPKeyBlock indexPkey Index		Guid added to partition
4	prtn_create	-PKey • Prtn Name		Partition created
5	prtn_delete	 PKey Delete Reason 	Delete Reason: 0 - empty prtn 1 - duplicate prtn 2 - sm shutdown	Partition deleted
6	port_virt_discover	Port GuidTop Index		Port virtualization discovered
7	vport_state_change	Port GuidVPort GuidVPort Index	VPort State: 1 - Down 2 - Init	Vport state changed

ACtivity ID	Activity Name	Additional Fields	Comments	Description
		VNode GuidVPort State	3 - ARMED 4 - Active	
8	mcg_tree_calc	mlid		MCast group tree calculated
9	routing_succeed	routing engine name		Routing done successfully
10	routing_failed	routing engine name		Routing failed
11	ucast_cache_invali- dated			ucast cache invalidated
12	ucast_cache_rout- ing_done			ucast cache routing done

Offsweep Balancing

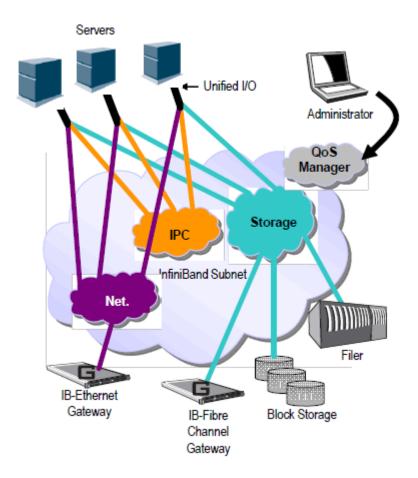
When working with minhop/dor/updn, subnet manager can re-balance routing during idle time (between sweeps).

- offsweep_balancing_enabled enables/disables the feature. Examples:
 - offsweep_balancing_enabled = TRUE
 - offsweep_balancing_enabled = FALSE (default)
- offsweep_balancing_window defines window of seconds to wait after sweep before starting the re-balance process. Applicable only if offsweep_balancing_enabled=TRUE. Example:

offsweep_balancing_window = 180 (default)

QoS - Quality of Service

Quality of Service (QoS) requirements stem from the realization of I/O consolidation over an IB network. As multiple applications and ULPs share the same fabric, a means is needed to control their use of network resources.



The basic need is to differentiate the service levels provided to different traffic flows, such that a policy can be enforced and can control each flow utilization of fabric resources.

The InfiniBand Architecture Specification defines several hardware features and management interfaces for supporting QoS:

Up to 15 Virtual Lanes (VL) carry traffic in a non-blocking manner

- Arbitration between traffic of different VLs is performed by a two-priority-level weighted round robin arbiter. The arbiter is programmable with a sequence of (VL, weight) pairs and a maximal number of high priority credits to be processed before low priority is served
- Packets carry class of service marking in the range 0 to 15 in their header SL field
- Each switch can map the incoming packet by its SL to a particular output VL, based on a programmable table VL=SL-to-VL-MAP(in-port, out-port, SL)
- The Subnet Administrator controls the parameters of each communication flow by providing them as a response to Path Record (PR) or MultiPathRecord (MPR) queries

DiffServ architecture (IETF RFC 2474 & 2475) is widely used in highly dynamic fabrics. The following subsections provide the functional definition of the various software elements that enable a DiffServ-like architecture over the NVIDIA OFED software stack.

QoS Architecture

QoS functionality is split between the SM/SA, CMA and the various ULPs. We take the "chronology approach" to describe how the overall system works.

- The network manager (human) provides a set of rules (policy) that define how the network is being configured and how its resources are split to different QoS-Levels. The policy also define how to decide which QoS-Level each application or ULP or service use.
- 2. The SM analyzes the provided policy to see if it is realizable and performs the necessary fab-ric setup. Part of this policy defines the default QoS-Level of each partition. The SA is enhanced to match the requested Source, Destination, QoS-Class, Service-ID, PKey against the policy, so clients (ULPs, programs) can obtain a policy enforced QoS. The SM may also set up partitions with appropriate IPoIB broadcast group. This broadcast group carries its QoS attributes: SL, MTU, RATE, and Packet Lifetime.
- 3. IPoIB is being setup. IPoIB uses the SL, MTU, RATE and Packet Lifetime available on the multicast group which forms the broadcast group of this partition.
- 4. MPI which provides non IB based connection management should be configured to run using hard coded SLs. It uses these SLs for every QP being opened.
- 5. ULPs that use CM interface (like SRP) have their own pre-assigned Service-ID and use it while obtaining PathRecord/MultiPathRecord (PR/MPR) for establishing connections. The SA receiving the PR/MPR matches it against the policy and returns the appropriate PR/MPR including SL, MTU, RATE and Lifetime.
- 6. ULPs and programs (e.g. SDP) use CMA to establish RC connection provide the CMA the target IP and port number. ULPs might also provide QoS-Class. The CMA then creates Service-ID for the ULP and passes this ID and optional QoS-Class in the PR/MPR request. The resulting PR/MPR is used for configuring the connection QP.

PathRecord and Multi Path Record Enhancement for QoS:

As mentioned above, the PathRecord and MultiPathRecord attributes are enhanced to carry the Service-ID which is a 64bit value. A new field QoS-Class is also provided.

A new capability bit describes the SM QoS support in the SA class port info. This approach provides an easy migration path for existing access layer and ULPs by not introducing new set of PR/MPR attributes.

Supported Policy

The QoS policy, which is specified in a stand-alone file, is divided into the following four subsections:

Port Group

A set of CAs, Routers or Switches that share the same settings. A port group might be a partition defined by the partition manager policy, list of GUIDs, or list of port names based on NodeDescription.

Fabric Setup

Defines how the SL2VL and VLArb tables should be set up.



In OFED this part of the policy is ignored. SL2VL and VLArb tables should be configured in the OpenSM options file (opensm.opts).

QoS-Levels Definition

This section defines the possible sets of parameters for QoS that a client might be mapped to. Each set holds SL and optionally: Max MTU, Max Rate, Packet Lifetime and Path Bits.



Matching Rules

A list of rules that match an incoming PR/MPR request to a QoS-Level. The rules are processed in order such as the first match is applied. Each rule is built out of a set of match expressions which should all match for the rule to apply. The matching expressions are defined for the following fields:

- SRC and DST to lists of port groups
- Service-ID to a list of Service-ID values or ranges
- QoS-Class to a list of QoS-Class values or ranges

CMA Features

The CMA interface supports Service-ID through the notion of port space as a prefix to the port number, which is part of the sockaddr provided to rdma_resolve_add(). The CMA also allows the ULP (like SDP) to propagate a request for a specific QoS-Class. The CMA uses the provided QoS-Class and Service-ID in the sent PR/MPR.

IPolB

IPoIB queries the SA for its broadcast group information and uses the SL, MTU, RATE and Packet Lifetime available on the multicast group which forms this broadcast group.

SRP

The current SRP implementation uses its own CM callbacks (not CMA). So SRP fills in the Service-ID in the PR/MPR by itself and use that information in setting up the QP.

SRP Service-ID is defined by the SRP target I/O Controller (it also complies with IBTA Service-ID rules). The Service-ID is reported by the I/O Controller in the ServiceEntries

DMA attribute and should be used in the PR/MPR if the SA reports its ability to handle QoS PR/MPRs.

IP over InfiniBand (IPoIB)

Upper Layer Protocol (ULP)

The IP over IB (IPoIB) ULP driver is a network interface implementation over InfiniBand. IPoIB encapsulates IP datagrams over an InfiniBand Datagram transport service. The IPoIB driver, ib_ipoib, exploits the following capabilities:

- VLAN simulation over an InfiniBand network via child interfaces
- High Availability via Bonding
- Varies MTU values:
 - up to 4k in Datagram mode
- Uses any ConnectX[®] IB ports (one or two)
- Inserts IP/UDP/TCP checksum on outgoing packets
- Calculates checksum on received packets
- Support net device TSO through ConnectX[®] LSO capability to defragment large data- grams to MTU quantas.

IPoIB also supports the following software based enhancements:

- Giant Receive Offload
- NAPI
- Ethtool support

Enhanced IPolB

Enhanced IPoIB feature enables offloading ULP basic capabilities to a lower vendor specific driver, in order to optimize IPoIB data path. This will allow IPoIB to support multiple

stateless offloads, such as RSS/TSS, and better utilize the features supported, enabling IPoIB datagram to reach peak performance in both bandwidth and latency.

Enhanced IPoIB supports/performs the following:

- Stateless offloads (RSS, TSS)
- Multi queues
- Interrupt moderation
- Multi partitions optimizations
- Sharing send/receive Work Queues
- Vendor specific optimizations
- UD mode only

Port Configuration

The physical port MTU (indicates the port capability) default value is 4k, whereas the IPoIB port MTU ("logical" MTU) default value is 2k as it is set by the OpenSM.

To change the IPoIB MTU to 4k, edit the OpenSM partition file in the section of IPoIB setting as follow:

Default=0xffff, ipoib, mtu=5 : ALL=full;

where:

"mtu=5" indicates that all IPoIB ports in the fabric are using 4k MTU, ("mtu=4" indi- cates 2k MTU)

IPoIB Configuration

Unless you have run the installation script mInxofedinstall with the flag '-n', then IPoIB has not been configured by the installation. The configuration of IPoIB requires assigning an IP address and a subnet mask to each HCA port, like any other network adapter card (i.e., you need to prepare a file called ifcfg-ib<n> for each port). The first port on the first HCA in the host is called interface ib0, the second port is called ib1, and so on. IPoIB configuration can be based on DHCP or on a static configuration that you need to supply (see below). You can also apply a manual configuration that persists only until the next reboot or driver restart (see below).

IPoIB Configuration Based on DHCP

Setting an IPoIB interface configuration based on DHCP is performed similarly to the configuration of Ethernet interfaces. In other words, you need to make sure that IPoIB configuration files include the following line:

• For RedHat:

BOOTPROTO=dhcp

• For SLES:

BOOTPROTO='dchp'

(i) Note

If IPoIB configuration files are included, ifcfg-ib<n> files will be installed under:

/etc/sysconfig/network-scripts/ on a RedHat machine

/etc/sysconfig/network/ on a SuSE machine.



A patch for DHCP may be required for supporting IPolB. For further information, please see the REAME file available under the docs/dhcp/ directory.

(i) Note

Red Hat Enterprise Linux 7 supports assigning static IP addresses to InfiniBand IPoIB interfaces. However, as these interfaces do not have a normal hardware Ethernet address, a different method of specifying a unique identifier for the IPoIB interface must be used. The standard is to use the option dhcpclient-identifier= construct to specify the IPoIB interface's dhcpclient-identifier field. The DHCP server host construct supports at most one hardware Ethernet and one dhcp-client-identifier entry per host stanza. However, there may be more than one fixed-address entry and the DHCP server will automatically respond with an address that is appropriate for the network that the DHCP request was received on.

Standard DHCP fields holding MAC addresses are not large enough to contain an IPoIB hardware address. To overcome this problem, DHCP over InfiniBand messages convey a client identifier field used to identify the DHCP session. This client identifier field can be used to associate an IP address with a client identifier value, such that the DHCP server will grant the same IP address to any client that conveys this client identifier.

The length of the client identifier field is not fixed in the specification. For the *NVIDIA OFED for Linux* package, it is recommended to have IPoIB use the same format that FlexBoot uses for this client identifier.

DHCP Server

In order for the DHCP server to provide configuration records for clients, an appropriate configuration file needs to be created. By default, the DHCP server looks for a configuration file called dhcpd.conf under /etc. You can either edit this file or create a new one and provide its full path to the DHCP server using the -cf flag (See a file example at docs/dhcpd.conf).

The DHCP server must run on a machine which has loaded the IPoIB module. To run the DHCP server from the command line, enter:

```
dhcpd <IB network interface name> -d
```

Example:

```
host1# dhcpd ib0 -d
```

DHCP Client (Optional)

(i) Note

A DHCP client can be used if you need to prepare a diskless machine with an IB driver.

In order to use a DHCP client identifier, you need to first create a configuration file that defines the DHCP client identifier.

Then run the DHCP client with this file using the following command:

dhclient -cf <client conf file> <IB network interface name>

Example of a configuration file for the ConnectX (PCI Device ID 26428), called dhclient.conf:

```
The value indicates a hexadecimal number interface "ib1" {
  send dhcp-client-identifier
  ff:00:00:00:00:00:02:00:00:02:c9:00:00:02:c9:03:00:00:10:39;
```

}

۶.

Example of a configuration file for InfiniHost III Ex (PCI Device ID 25218), called dhclient.conf:

```
The value indicates a hexadecimal number interface "ib1" {
  send dhcp-client-identifier
  20:00:55:04:01:fe:80:00:00:00:00:00:00:00:02:c9:02:00:23:13:92;
}
```

In order to use the configuration file, run

```
host1# dhclient -cf dhclient.conf ib1
```

Static IPoIB Configuration

If you wish to use an IPoIB configuration that is not based on DHCP, you need to supply the installation script with a configuration file (using the '-n' option) containing the full IP configuration. The IPoIB configuration file can specify either or both of the following data for an IPoIB interface:

- A static IPoIB configuration
- An IPoIB configuration based on an Ethernet configuration

See your Linux distribution documentation for additional information about configuring IP addresses.

The following code lines are an excerpt from a sample IPoIB configuration file:

```
# Static settings; all values provided by this file
IPADDR_ib0=10.4.3.175
```

```
NETMASK_ib0=255.255.0.0
NETWORK_ib0=10.4.0.0
BROADCAST_ib0=10.4.255.255
ONBOOT_ib0=1
# Based on eth0; each '*' will be replaced with a corresponding
octet
# from eth0.
LAN_INTERFACE_ib0=eth0
IPADDR_ib0=10.4.'*'.'*'
NETMASK_ib0=255.255.0.0
NETWORK_ib0=10.4.0.0
BROADCAST_ib0=10.4.255.255
ONBOOT_ib0=1
# Based on the first eth<n> interface that is found (for n=0,1,...);
# each '*' will be replaced with a corresponding octet from
eth<n>.
LAN_INTERFACE_ib0=
IPADDR_ib0=10.4.'*'.'*'
NETMASK_ib0=255.255.0.0
NETWORK_ib0=10.4.0.0
BROADCAST_ib0=10.4.255.255
ONBOOT_ib0=1
```

Manually Configuring IPolB

(i) Note

This manual configuration persists only until the next reboot or driver restart.

> To manually configure IPoIB for the default IB partition (VLAN), perform the following steps:

- 1. Configure the interface by entering the ifconfig command with the following items:
 - The appropriate IB interface (ib0, ib1, etc.)
 - The IP address that you want to assign to the interface
 - The netmask keyword
 - The subnet mask that you want to assign to the interface

The following example shows how to configure an IB interface:

host1\$ ifconfig ib0 10.4.3.175 netmask 255.255.0.0

2. (Optional) Verify the configuration by entering the ifconfig command with the appropriate interface identifier *ib#* argument.

The following example shows how to verify the configuration:

```
host1$ ifconfig ib0
b0 Link encap:UNSPEC HWaddr 80-00-04-04-FE-80-00-00-00-00-00-
00-00-00-00
inet addr:10.4.3.175 Bcast:10.4.255.255 Mask:255.255.0.0
UP BROADCAST MULTICAST MTU:65520 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:128
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

3. Repeat the first two steps on the remaining interface(s).

Sub-interfaces

You can create sub-interfaces for a primary IPoIB interface to provide traffic isolation. Each such sub-interface (also called a child interface) has a different IP and network addresses from the primary (parent) interface. The default Partition Key (PKey), ff:ff, applies to the primary (parent) interface.

This section describes how to:

- Create a subinterface
- Remove a subinterface

Creating a Subinterface

In the following procedure, ib0 is used as an example of an IB sub-interface.

To create a child interface (sub-interface), follow this procedure:

≻

- 1. Decide on the PKey to be used in the subnet (valid values can be 0 or any 16-bit unsigned value). The actual PKey used is a 16-bit number with the most significant bit set. For example, a value of 1 will give a PKey with the value 0x8001.
- 2. Create a child interface by running:

host1\$ echo <PKey> > /sys/class/net/<IB
subinterface>/create_child

Example:

host1\$ echo 1 > /sys/class/net/ib0/create_child

This will create the interface ib0.8001.

3. Verify the configuration of this interface by running:

```
host1$ ifconfig <subinterface>.<subinterface PKey>
```

Using the example of the previous step:

```
host1$ ifconfig ib0.8001
ib0.8001 Link encap:UNSPEC HWaddr 80-00-00-4A-FE-80-00-00-
00-00-00-00-00-00
BROADCAST MULTICAST MTU:2044 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:128
RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

- 4. As can be seen, the interface does not have IP or network addresses. To configure those, you should follow the manual configuration procedure described in "<u>Manually</u> <u>Configuring IPolB</u>" section above.
- 5. To be able to use this interface, a configuration of the Subnet Manager is needed so that the PKey chosen, which defines a broadcast address, be recognized.

Removing a Subinterface

To remove a child interface (subinterface), run:

```
≻
```

```
echo <subinterface PKey> /sys/class/net/<ib_interface>/delete_child
```

Using the example of \underline{t} he second step from the previous chapter:

```
echo 0x8001 > /sys/class/net/ib0/delete_child
```

Note that when deleting the interface you must use the PKey value with the most significant bit set (e.g., 0x8000 in the example above).

Verifying IPoIB Functionality

To verify your configuration and IPoIB functionality are successful, perform the following steps:

1. Verify the IPoIB functionality by using the ifconfig command.

The following example shows how two IB nodes are used to verify IPoIB functionality. In the following example, IB node 1 is at 10.4.3.175, and IB node 2 is at 10.4.3.176:

host1# ifconfig ib0 10.4.3.175 netmask 255.255.0.0 host2# ifconfig ib0 10.4.3.176 netmask 255.255.0.0

- 2. Enter the ping command from 10.4.3.175 to 10.4.3.176.
- 3. The following example shows how to enter the ping command:

```
host1# ping -c 5 10.4.3.176
PING 10.4.3.176 (10.4.3.176) 56(84) bytes of data.
64 bytes from 10.4.3.176: icmp_seq=0 ttl=64 time=0.079 ms
64 bytes from 10.4.3.176: icmp_seq=1 ttl=64 time=0.044 ms
64 bytes from 10.4.3.176: icmp_seq=2 ttl=64 time=0.055 ms
64 bytes from 10.4.3.176: icmp_seq=3 ttl=64 time=0.049 ms
64 bytes from 10.4.3.176: icmp_seq=4 ttl=64 time=0.065 ms
--- 10.4.3.176 ping statistics ---
```

```
5 packets transmitted, 5 received, 0% packet loss, time
3999ms rtt min/avg/max/mdev = 0.044/0.058/0.079/0.014 ms, pipe 2
```

Bonding IPolB

To create an interface configuration script for the ibX and bondX interfaces, you should use the standard syntax (depending on your OS).

Bonding of IPoIB interfaces is accomplished in the same manner as would bonding of Ethernet interfaces: via the Linux Bonding Driver.

- Network Script files for IPoIB slaves are named after the IPoIB interfaces (e.g: ifcfg-ib0)
- The only meaningful bonding policy in IPoIB is High-Availability (bonding mode number 1, or active-backup)
- Bonding parameter "fail_over_mac" is meaningless in IPoIB interfaces, hence, the only supported value is the default: 0

For a persistent bonding IPoIB Network configuration, use the same Linux Network Scripts semantics, with the following exceptions/ additions:

• In the bonding master configuration file (e.g: ifcfg-bond0), in addition to Linux bonding semantics, use the following parameter: MTU=65520

COND

(i) Note

For IPoIB slaves, use MTU=2044. If you do not set the correct MTU or do not set MTU at all, performance of the interface might decrease.

Dynamically Connected Transport (DCT)

• In the bonding slave configuration file (e.g: ifcfg-ib0), use the same Linux Network Scripts semantics. In particular: DEVICE=ib0

- In the bonding slave configuration file (e.g: ifcfg-ib0.8003), the line TYPE=InfiniBand is necessary when using bonding over devices configured with partitions (p_key)
- For RHEL users:

In /etc/modprobe.b/bond.conf add the following lines:

alias bond0 bonding

• For SLES users:

It is necessary to update the MANDATORY_DEVICES environment variable in /etc/sysconfig/network/config with the names of the IPoIB slave devices (e.g. ib0, ib1, etc.). Otherwise, bonding master may be created before IPoIB slave interfaces at boot time.

It is possible to have multiple IPoIB bonding masters and a mix of IPoIB bonding master and Ethernet bonding master. However, It is NOT possible to mix Ethernet and IPoIB slaves under the same bonding master.

) Note

Restarting openibd does no keep the bonding configuration via Network Scripts. You have to restart the network service in order to bring up the bonding master. After the configuration is saved, restart the network service by running: /etc/init.d/network restart.

Dynamic PKey Change

Dynamic PKey change means the PKey can be changed (add/removed) in the SM database and the interface that is attached to that PKey is updated immediately without the need to restart the driver.

If the PKey is already configured in the port by the SM, the child-interface can be used immediately. If not, the interface will be ready to use only when SM adds the relevant PKey value to the port after the creation of the child interface. No additional configuration is required once the child-interface is created.

Precision Time Protocol (PTP) over IPolB

This feature allows for accurate synchronization between the distributed entities over the network. The synchronization is based on symmetric Round Trip Time (RTT) between the master and slave devices.

This feature is enabled by default, and is also supported over PKey interfaces.

For more on the PTP feature, refer to <u>Running Linux PTP with ConnectX-4/ConnectX-5/ConnectX-6</u> Community post.

For further information on Time-Stamping, follow the steps in "Time-Stamping Service".

One Pulse Per Second (1PPS) over IPoIB

1PPS is a time synchronization feature that allows the adapter to be able to send or receive 1 pulse per second on a dedicated pin on the adapter card using an SMA connector (SubMiniature version A). Only one pin is supported and could be configured as 1PPS in or 1PPS out.

For further information, refer to HowTo Test 1PPS on NVIDIA Adapters Community post.

Advanced Transport

Atomic Operations

Atomic Operations in mlx5 Driver

To enable atomic operation with this endianness contradiction, use the <code>ibv_create_qp</code> to create the QP and set the <code>IBV_QP_CREATE_ATOMIC_BE_REPLY</code> flag on <code>create_flags</code>.

XRC - eXtended Reliable Connected Transport Service for InfiniBand

XRC allows significant savings in the number of QPs and the associated memory resources required to establish all to all process connectivity in large clusters.

It significantly improves the scalability of the solution for large clusters of multicore endnodes by reducing the required resources.

For further details, please refer to the "Annex A14 Supplement to InfiniBand Architecture Specification Volume 1.2.1"

A new API can be used by user space applications to work with the XRC transport. The legacy API is currently supported in both binary and source modes, however it is deprecated. Thus we recommend using the new API.

The new verbs to be used are:

- ibv_open_xrcd/ibv_close_xrcd
- ibv_create_srq_ex
- ibv_get_srq_num
- ibv_create_qp_ex
- ibv_open_qp

Please use ibv_xsrq_pingpong for basic tests and code reference. For detailed information regarding the various options for these verbs, please refer to their appropriate man pages.

Dynamically Connected Transport (DCT)

Dynamically Connected transport (DCT) service is an extension to transport services to enable a higher degree of scalability while maintaining high performance for sparse traffic. Utilization of DCT reduces the total number of QPs required system wide by having Reliable type QPs dynamically connect and disconnect from any remote node. DCT connections only stay connected while they are active. This results in smaller memory footprint, less overhead to set connections and higher on-chip cache utilization and hence increased performance. DCT is supported only in mlx5 driver.



Please note that ConnectX-4 supports DCT v0 and ConnectX-5 and above support DCT v1. DCTv0 and DCT v1 are not interoperable.

MPI Tag Matching and Rendezvous Offloads

(i) Note

Supported in ConnectX®-5 and above adapter cards.

Tag Matching and Rendezvous Offloads is a technology employed by NVIDIA to offload the processing of MPI messages from the host machine onto the network card. Employing this technology enables a zero copy of MPI messages, i.e. messages are scattered directly to the user's buffer without intermediate buffering and copies. It also provides a complete rendezvous progress by NVIDIA devices. Such overlap capability enables the CPU to perform the application's computational tasks while the remote data is gathered by the adapter.

For more information Tag Matching Offload, please refer to the <u>Understanding MPI Tag</u> <u>Matching and Rendezvous Offloads (ConnectX-5)</u> Community post .

Optimized Memory Access

Memory Region Re-registration

Memory Region Re-registration allows the user to change attributes of the memory region. The user may change the PD, access flags or the address and length of the memory region. Memory

region supports contagious pages allocation. Consequently, it de-registers memory region followed by register memory region. Where possible, resources are reused instead of de-allocated and reallocated.

Example:

```
int ibv_rereg_mr(struct ibv_mr *mr, int flags, struct ibv_pd *pd,
void *addr, size_t length, uint64_t access, struct
ibv_rereg_mr_attr *attr);
```

@mr:	The memory region to modify.	
@flags:	A bit-mask used to indicate which of the following properties of the memory region are being modified. Flags should be one of: IBV_REREG_MR_CHANGE_TRANSLATION /* Change translation (location and length) */ IBV_REREG_MR_CHANGE_PD/* Change protection domain*/ IBV_REREG_MR_CHANGE_ACCESS/* Change access flags*/	
@pd:	If IBV_REREG_MR_CHANGE_PD is set in flags, this field specifies the new protection domain to associated with the memory region, otherwise, this parameter is ignored.	
@addr:	If IBV_REREG_MR_CHANGE_TRANSLATION is set in flags, this field specifies the start of the virtual address to use in the new translation, otherwise, this parameter is ignored.	
@length:	If IBV_REREG_MR_CHANGE_TRANSLATION is set in flags, this field specifies the length of the virtual address to use in the new translation, otherwise, this parameter is ignored.	
@access:	If IBV_REREG_MR_CHANGE_ACCESS is set in flags, this field specifies the new memory access rights, otherwise, this parameter is ignored. Could be one of the following: IBV_ACCESS_LOCAL_WRITE IBV_ACCESS_REMOTE_WRITE IBV_ACCESS_REMOTE_READ IBV_ACCESS_ALLOCATE_MR /* Let the library allocate the memory for * the user, tries to get contiguous pages */	
@attr:	Future extensions	

ibv_rereg_mr returns 0 on success, or the value of an errno on failure (which indicates the error reason). In case of an error, the MR is in undefined state. The user needs to call ibv_dereg_mr in order to release it.

Please note that if the MR (Memory Region) is created as a Shared MR and a translation is requested, after the call, the MR is no longer a shared MR. Moreover, Re-registration of MRs that uses NVIDIA PeerDirect[™] technology are not supported.

Memory Window

Memory Window allows the application to have a more flexible control over remote access to its memory. It is available only on physical functions/native machines The two types of Memory Windows supported are: type 1 and type 2B.

Memory Windows are intended for situations where the application wants to:

- Grant and revoke remote access rights to a registered region in a dynamic fashion with less of a performance penalty
- Grant different remote access rights to different remote agents and/or grant those rights over different ranges within registered region

For further information, please refer to the InfiniBand specification document.

Memory Windows API cannot co-work with peer memory clients (PeerDirect).

Query Capabilities

Note

Memory Windows are available if and only the hardware supports it. To verify whether Memory Windows are available, run ibv_query_device.

For example:

Memory Window Allocation

Allocating memory window is done by calling the <code>ibv_alloc_mw</code> verb.

```
type_mw = IBV_MW_TYPE_2/ IBV_MW_TYPE_1
mw = ibv_alloc_mw(pd, type_mw);
```

Binding Memory Windows

After being allocated, memory window should be bound to a registered memory region. Memory Region should have been registered using the IBV_ACCESS_MW_BIND access flag.

For further information on how to bind memory windows, please see <u>rdma-core man</u> <u>page</u>.

Invalidating Memory Window

Before rebinding Memory Window type 2, it must be invalidated using ibv_post_send - see <u>here</u>.

Deallocating Memory Window

Deallocating memory window is done using the ibv_dealloc_mw verb.

ibv_dealloc_mw(mw);

User-Mode Memory Registration (UMR)

User-mode Memory Registration (UMR) is a fast registration mode which uses send queue. The UMR support enables the usage of RDMA operations and scatters the data at the remote side through the definition of appropriate memory keys on the remote side.

UMR enables the user to:

- Create indirect memory keys from previously registered memory regions, including creation of KLM's from previous KLM's. There are not data alignment or length restrictions associated with the memory regions used to define the new KLM's.
- Create memory regions, which support the definition of regular non-contiguous memory regions.

On-Demand-Paging (ODP)

On-Demand-Paging (ODP) is a technique to alleviate much of the shortcomings of memory registration. Applications no longer need to pin down the underlying physical pages of the address space, and track the validity of the mappings. Rather, the HCA requests the latest translations from the OS when pages are not present, and the OS invalidates translations which are no longer valid due to either non-present pages or mapping changes. ODP does not support contiguous pages.

ODP can be further divided into 2 subclasses: Explicit and Implicit ODP.

• Explicit ODP

In Explicit ODP, applications still register memory buffers for communication, but this operation is used to define access control for IO rather than pin-down the pages. ODP Memory Region (MR) does not need to have valid mappings at registration time.

• Implicit ODP

In Implicit ODP, applications are provided with a special memory key that represents their complete address space. This all IO accesses referencing this key (subject to the access rights associated with the key) does not need to register any virtual address range.

Query Capabilities

On-Demand Paging is available if both the hardware and the kernel support it. To verify whether ODP is supported, run ibv_query_device.

For further information, please refer to the *ibv query device manual page*.

Registering ODP Explicit and Implicit MR

ODP Explicit MR is registered after allocating the necessary resources (e.g. PD, buffer), while ODP implicit MR registration provides an implicit lkey that represents the complete address space.

For further information, please refer to the *ibv_reg_mr_manual page*.

De-registering ODP MR

ODP MR is deregistered the same way a regular MR is deregistered:

```
ibv_dereg_mr(mr);
```

Advice MR Verb

The driver can pre-fetch a given range of pages and map them for access from the HCA. The advice MR verb is applicable for ODP MRs only.

For further information, please refer to the *ibv_advise_mr manual page*.

ODP Statistics

To aid in debugging and performance measurements and tuning, ODP support includes an extensive set of statistics.

For further information, please refer to rdma-statistics manual page.

Inline-Receive

The HCA may write received data to the Receive CQE. Inline-Receive saves PCIe Read transaction since the HCA does not need to read the scatter list. Therefore, it improves performance in case of short receive-messages.

On poll CQ, the driver copies the received data from CQE to the user's buffers.

Inline-Receive is enabled by default and is transparent to the user application. To disable it globally, set MLX5_SCATTER_TO_CQE environment variable to the value of 0. Otherwise, disable it on a specific QP using mlx5dv_create_qp() with MLX5DV_QP_CREATE_DISABLE_SCATTER_TO_CQE.

For further information, please refer to the manual page of mlx5dv_create_qp().

NVIDIA PeerDirect

NVIDIA PeerDirect[™] uses an API between IB CORE and peer memory clients, (e.g. GPU cards) to provide access to an HCA to read/write peer memory for data buffers. As a result, it allows RDMA-based (over InfiniBand/RoCE) application to use peer device computing power, and RDMA interconnect at the same time without copying the data between the P2P devices.

For example, PeerDirect is being used for GPUDirect RDMA.

Detailed description for that API exists under MLNX OFED installation, please see docs/readme_and_user_manual/PEER_MEMORY_API.txt.

PeerDirect Async

Mellanox PeerDirect Async sub-system gives PeerDirect hardware devices, such as GPU cards, dedicated AS accelerators, and so on, the ability to take control over HCA in critical path offloading CPU. To achieve this, there is a set of verb calls and structures providing application with abstract description of operation sequences intended to be executed by peer device.

Relaxed Ordering (RSYNC)

j Note

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

In GPU systems with relaxed ordering, RSYNC callback will be invoked to ensure memory consistency. The registration and implementation of the callback will be done using an

external module provided by the system vendor. Loading the module will register the callback in MLNX_OFED to be used later to guarantee memory operations order.

CPU Overhead Distribution

When creating a CQ using the ibv_create_cq() API, a "comp_vector" argument is sent. If the value set for this argument is 0, while the CPU core executing this verb is not equal to zero, the driver assigns a completion EQ with the least CQs reporting to it. This method is used to distribute CQs amongst available completions EQ. To assign a CQ to a specific EQ, the EQ needs to be specified in the comp_vector argument.

Out-of-Order (OOO) Data Placement

(i) Note

This feature is only supported on:

- ConnectX-5 adapter cards and above
- RC and XRC QPs
- DC transport

Overview

In certain fabric configurations, InfiniBand packets for a given QP may take up different paths in a network from source to destination. This results into packets being received in an out-of-order manner. These packets can now be handled instead of being dropped, in order to avoid retransmission, by:

- Achieving better network utilization
- Decreasing latency

Data will be placed into host memory in an out-of-order manner when out-of-order messages are received.

For information on how to set up out-of-order processing by the QP, please refer to <u>HowTo Configure Adaptive Routing and SHIELD</u> Community post.

IB Router

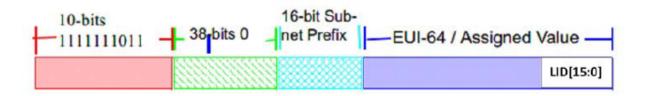
IB router provides the ability to send traffic between two or more IB subnets thereby potentially expanding the size of the network to over 40k end-ports, enabling separation and fault resilience between islands and IB subnets, and enabling connection to different topologies used by different subnets.

The forwarding between the IB subnets is performed using GRH lookup. The IB router's basic functionality includes:

- Removal of current L2 LRH (local routing header)
- Routing
- table lookup using GID from GRH
- Building new LRH according to the destination according to the routing table

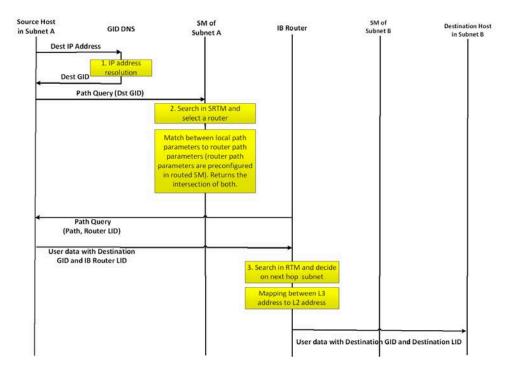
The DLID in the new LRH is built using simplified GID-to-LID mapping (where LID = 16 LSB bits of GID) thereby not requiring to send for ARP query/lookup.

Local Unicast GID Format



For this to work, the SM allocates an alias GID for each host in the fabric where the alias GID = {subnet prefix[127:64], reserved[63:16], LID[15:0}. Hosts should use alias GIDs in order to transmit traffic to peers on remote subnets.

Host-to-Host IB Router Unicast Flow



- For information on the architecture and functionality of IB Router, refer to <u>IB Router</u> <u>Architecture and Functionality</u> Community post.
- For information on IB Router configuration, refer to <u>HowTo Configure IB Routers</u> Community post.

Storage Protocols

There are several storage protocols that use the advantage of InfiniBand and RDMA for performance reasons (high throughput, low latency and low CPU utilization). In this chapter we will discuss the following protocols:

- **SCSI RDMA Protocol (SRP)** is designed to take full advantage of the protocol offload and RDMA features provided by the InfiniBand architecture.
- **iSCSI Extensions for RDMA (iSER)** is an extension of the data transfer model of iSCSI, a storage networking standard for TCP/IP. It uses the iSCSI components while taking the advantage of the RDMA protocol suite. ISER is implemented on various stor- age targets such as TGT, LIO, SCST and out of scope of this manual.

For various ISER targets configuration steps, troubleshooting and debugging, as well as other implementation of storage protocols over RDMA (such as Ceph over RDMA, nbdX and more) refer to Storage Solutions on the Community website.

- **Lustre** is an open-source, parallel distributed file system, generally used for largescale cluster computing that supports many requirements of leadership class HPC simulation environments.
- NVM Express[™] over Fabrics (NVME-oF)
 - NVME-oF is a technology specification for networking storage designed to enable NVMe message-based commands to transfer data between a host computer and a target solid-state storage device or system over a network such as Ethernet, Fibre Channel, and InfiniBand. Tunneling NVMe commands through an RDMA fabric provides a high throughput and a low latency. This is an alternative to the SCSi based storage networking protocols.
 - NVME-oF Target Offload is an implementation of the new NVME-oF standard Target (server) side in hardware. Starting from ConnectX-5 family cards, all regular IO requests can be processed by the HCA, with the HCA sending IO requests directly to a real NVMe PCI device, using peer-to-peer PCI communications. This means that excluding connection management and error flows, no CPU utilization will be observed during NVME-oF traffic.

For further information, please refer to Storage Solutions on the Community website (<u>enterprise-support.nvidia.com/s/</u>).

SRP - SCSI RDMA Protocol

The 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 Initiator controls the connection to an SRP Target in order to provide access to remote storage devices across an InfiniBand fabric. The kSRP Target resides in an IO unit and provides storage services.

SRP Initiator

This SRP Initiator is based on open source from OpenFabrics (<u>www.openfabrics.org</u>) that implements the SCSI RDMA Protocol-2 (SRP-2). SRP-2 is described in Document # T10/1524-D available from <u>http://www.t10.org</u>.

The SRP Initiator supports

• Basic SCSI Primary Commands -3 (SPC-3)

(www.t10.org/ftp/t10/drafts/spc3/spc3r21b.pdf)

• Basic SCSI Block Commands -2 (SBC-2)

(www.t10.org/ftp/t10/drafts/sbc2/sbc2r16.pdf)

• Basic functionality, task management and limited error handling

j Note

This package, however, does not include an SRP Target.

Loading SRP Initiator

To load the SRP module either:

• Execute the modprobe ib_srp command after the OFED driver is up.

or

- 1. Change the value of SRP_LOAD in /etc/infiniband/openib.conf to "yes".
- 2. Run /etc/init.d/openibd restart for the changes to take effect.

i) Note

When loading the ib_srp module, it is possible to set the module parameter srp_sg_tablesize. This is the maximum number of gather/scatter entries per I/O (default: 12).

SRP Module Parameters

When loading the SRP module, the following parameters can be set (viewable by the "modinfo ib_srp" command):

cmd_sg_entri es	Default number of gather/scatter entries in the SRP command (default is 12, max 255)					
allow_ext_sg	Default behavior when there are more than cmd_sg_entries S/G entries after mapping; fails the request when false (default false)					
topspin_work arounds	Enable workarounds for Topspin/Cisco SRP target bugs					
reconnect_del ay	Time between successive reconnect attempts. Time between successive reconnect attempts of SRP initiator to a disconnected target until dev_loss_tmo timer expires (if enabled), after that the SCSI target will be removed					
fast_io_fail_tm o	Number of seconds between the observation of a transport layer error and failing all I/O. Increasing this timeout allows more tolerance to transport errors, however, doing so increases the total failover time in case of serious transport failure. Note: fast_io_fail_tmo value must be smaller than the value of recon- nect_delay					
dev_loss_tmo	Maximum number of seconds that the SRP transport should insulate transport layer errors. After this time has been exceeded the SCSI target is removed. Normally it is advised to set this to -1 (disabled) which will never remove the scsi_host. In deployments where different SRP targets are connected and disconnected frequently, it may be required to enable this timeout in order to clean old scsi_hosts representing targets that no longer exists					

Constraints between parameters:

- dev_loss_tmo, fast_io_fail_tmo, reconnect_delay cannot be all disabled or negative values.
- reconnect_delay must be positive number.
- fast_io_fail_tmo must be smaller than SCSI block device timeout.
- fast_io_fail_tmo must be smaller than dev_loss_tmo.

SRP Remote Ports Parameters

Several SRP remote ports parameters are modifiable online on existing connection.

echo 600 > /sys/class/srp_remote_ports/port-xxx/dev_loss_tmo

To modify fast_io_fail_tmo to 15 seconds:

echo 15 > /sys/class/srp_remote_ports/port-xxx/fast_io_fail_tmo

To modify reconnect_delay to 10 seconds:

echo 20 > /sys/class/srp_remote_ports/port-xxx/reconnect_delay

Manually Establishing an SRP Connection

The following steps describe how to manually load an SRP connection between the Initiator and an SRP Target. "<u>Automatic Discovery and Connection to Targets</u>" section explains how to do this automatically.

- Make sure that the ib_srp module is loaded, the SRP Initiator is reachable by the SRP Target, and that an SM is running.
- To establish a connection with an SRP Target and create an SRP (SCSI) device for that target under /dev, use the following command:

```
echo -n id_ext=[GUID value],ioc_guid=[GUID value],dgid=[port
GID value],\
pkey=ffff,service_id=[service[0] value] > \
/sys/class/infiniband_srp/srp-mlx[hca number]-[port
number]/add_target
```

See "<u>SRP Tools - ibsrpdm, srp_daemon and srpd Service Script</u>" section for instructions on how the parameters in this echo command may be obtained.

Notes:

- Execution of the above "echo" command may take some time
- The SM must be running while the command executes
- It is possible to include additional parameters in the echo command:
 - max_cmd_per_lun Default: 62
 - max_sect (short for max_sectors) sets the request size of a command
 - io_class Default: 0x100 as in rev 16A of the specification (In rev 10 the default was 0xff00)
 - tl_retry_count a number in the range 2..7 specifying the IB RC retry count. Default: 2
 - comp_vector, a number in the range 0..n-1 specifying the MSI-X completion vector. Some HCA's allocate multiple (n) MSI-X vectors per HCA port. If the IRQ affinity masks of these interrupts have been configured such that each MSI-X interrupt is handled by a different CPU then the comp_vector parameter can be used to spread the SRP completion workload over multiple CPU's.
 - cmd_sg_entries, a number in the range 1..255 that specifies the maximum number of data buffer descriptors stored in the SRP_CMD information unit itself. With allow_ext_sg=0 the parameter cmd_sg_entries defines the maximum S/G list length for a single SRP_CMD, and commands whose S/G list length exceeds this limit after S/G list collapsing will fail.
 - initiator_ext see "<u>Multiple Connections from Initiator InfiniBand Port to the</u> <u>Target</u>" section.
- To list the new SCSI devices that have been added by the echo command, you may use either of the following two methods:
 - Execute "fdisk -I". This command lists all devices; the new devices are included in this listing.
 - Execute "dmesg" or look at /var/log/messages to find messages with the names of the new devices.

SRP sysfs Parameters

Interface for making ib_srp connect to a new target. One can request ib_srp to connect to a new target by writing a comma-separated list of login parameters to this sysfs attribute. The supported parameters are:

id_ext	A 16-digit hexadecimal number specifying the eight byte identifier extension in the 16-byte SRP target port identifier. The target port identifier is sent by ib_srp to the target in the SRP_LOGIN_REQ request.			
ioc_guid	A 16-digit hexadecimal number specifying the eight byte I/O controller GUIE portion of the 16-byte target port identifier.			
dgid	A 32-digit hexadecimal number specifying the destination GID.			
pkey	A four-digit hexadecimal number specifying the InfiniBand partition key.			
service_id	A 16-digit hexadecimal number specifying the InfiniBand service ID used to establish communication with the SRP target. How to find out the value of the service ID is specified in the documentation of the SRP target.			
max_sect	A decimal number specifying the maximum number of 512-byte sectors to be transferred via a single SCSI command.			
max_cmd _per_lun	A decimal number specifying the maximum number of outstanding commands for a single LUN.			
io_class	A hexadecimal number specifying the SRP I/O class. Must be either 0xff00 (rev 10) or 0x0100 (rev 16a). The I/O class defines the format of the SRP initiator and target port identifiers.			
initiator_e xt	A 16-digit hexadecimal number specifying the identifier extension portion of the SRP initiator port identifier. This data is sent by the initiator to the target in the SRP_LOGIN_REQ request.			
cmd_sg_e ntries	$=$ 3^{-1} Lallow ext solution because the maximum S/(1)			
allow_ext_ sg	Whether ib_srp is allowed to include a partial memory descriptor list in an SRP_CMD instead of the entire list. If a partial memory descriptor list has been included in an SRP_CMD the remaining memory descriptors are communicated from initiator to target via an additional RDMA transfer. Setting allow_ext_sg to 1 increases the maximum amount of data that can			

	be transferred between initiator and target via a single SCSI command. Since not all SRP target implementations support partial memory descriptor lists the default value for this option is 0.
sg_tablesi ze	A number in the range 12048 specifying the maximum S/G list length the SCSI layer is allowed to pass to ib_srp. Specifying a value that exceeds cmd_sg_entries is only safe with partial memory descriptor list support enabled (allow_ext_sg=1).
comp_vec tor	A number in the range 0n-1 specifying the MSI-X completion vector. Some HCA's allocate multiple (n) MSI-X vectors per HCA port. If the IRQ affinity masks of these interrupts have been configured such that each MSI-X interrupt is handled by a different CPU then the comp_vector parameter can be used to spread the SRP completion workload over multiple CPU's.
tl_retry_co unt	A number in the range 27 specifying the IB RC retry count.

SRP Tools - ibsrpdm, srp_daemon and srpd Service Script

The OFED distribution provides two utilities: ibsrpdm and srp_daemon:

- They detect targets on the fabric reachable by the Initiator (Step 1)
- Output target attributes in a format suitable for use in the above "echo" command (Step 2)
- A service script srpd which may be started at stack startup

The utilities can be found under /usr/sbin/, and are part of the srptools RPM that may be installed using the OFED installation. Detailed information regarding the various options for these utilities are provided by their man pages.

Below, several usage scenarios for these utilities are presented.

ibsrpdm

ibsrpdm has the following tasks:

- 1. Detecting reachable targets.
 - 1. To detect all targets reachable by the SRP initiator via the default umad device (/sys/class/infiniband_mad/umad0), execute the following command:

This command will result into readable output information on each SRP Target detected. Sample:

```
IO Unit Info:
                     0103
          port LID:
          port GID:
fe80000000000000002c90200402bd5
          change ID:
                     0002
         max controllers: 0x10
       controller [ 1]
          GUID: 0002c90200402bd4
          vendor ID: 0002c9
          device ID: 005a44
          IO class : 0100
          ID:
                     LSI Storage Systems SRP Driver
200400a0b81146a1
           service entries: 1
          service[ 0]: 200400a0b81146a1 /
SRP.T10:200400A0B81146A1
```

2. To detect all the SRP Targets reachable by the SRP Initiator via another umad device, use the following command:

ibsrpdm -d <umad device>

- 2. Assisting in SRP connection creation.
 - To generate an output suitable for utilization in the "echo" command in <u>"Manually Establishing an SRP Connection</u>" section, add the '-c' option to ibsrpdm:

```
ibsrpdm -c
```

Sample output:

```
id_ext=200400A0B81146A1,ioc_guid=0002c90200402bd4,
dgid=fe800000000000000002c90200402bd5,pkey=ffff,service_id
```

2. To establish a connection with an SRP Target using the output from the 'ibsrpdm -c' example above, execute the following command:

```
echo -n
id_ext=200400A0B81146A1,ioc_guid=0002c90200402bd4,
dgid=fe800000000000000002c90200402bd5,pkey=ffff,service_id
> /sys/
class/infiniband_srp/srp-mlx5_0-1/add_target
```

The SRP connection should now be up; the newly created SCSI devices should appear in the listing obtained from the 'fdisk -I' command.

3. Discover reachable SRP Targets given an InfiniBand HCA name and port, rather than by just running /sys/class/infiniband_mad/umad<N> where <N> is a digit.

srpd

The srpd service script allows automatic activation and termination of the srp_daemon utility on all system live InfiniBand ports.

srp_daemon

srp_daemon utility is based on ibsrpdm and extends its functionality. In addition to the ibsrpdm functionality described above, srp_daemon can:

- Establish an SRP connection by itself (without the need to issue the "echo" command described in "<u>Manually Establishing an SRP Connection</u>" section)
- Continue running in background, detecting new targets and establishing SRP connections with them (daemon mode)
- Discover reachable SRP Targets given an infiniband HCA name and port, rather than just by /dev/umad<N> where <N> is a digit
- Enable High Availability operation (together with Device-Mapper Multipath)
- Have a configuration file that determines the targets to connect to:
- 1. srp_daemon commands equivalent to ibsrpdm:

```
"srp_daemon -a -o" is equivalent to "ibsrpdm"
"srp_daemon -c -a -o" is equivalent to "ibsrpdm -c"
```

Note: These srp_daemon commands can behave differently than the equivalent ibsrpdm command when /etc/srp_daemon.conf is not empty.

- 2. srp_daemon extensions to ibsrpdm.
- To discover SRP Targets reachable from the HCA device <InfiniBand HCA name> and the port <port num>, (and to generate output suitable for 'echo'), execute:

```
host1# srp_daemon -c -a -o -i <InfiniBand HCA name> -p <port
number>
```

Note: To obtain the list of InfiniBand HCA device names, you can either use the ibstat tool or run 'ls /sys/class/infiniband'.

- To both discover the SRP Targets and establish connections with them, just add the -e option to the above command.
- Executing srp_daemon over a port without the -a option will only display the reachable targets via the port and to which the initiator is not connected. If executing with the -e option it is better to omit -a.

- It is recommended to use the -n option. This option adds the initiator_ext to the connecting string (see "<u>Multiple Connections from Initiator InfiniBand Port to the</u> <u>Target</u>" section).
- srp_daemon has a configuration file that can be set, where the default is /etc/srp_daemon.conf. Use the -f to supply a different configuration file that configures the targets srp_daemon is allowed to connect to. The configuration file can also be used to set values for additional parameters (e.g., max_cmd_per_lun, max_sect).
- A continuous background (daemon) operation, providing an automatic ongoing detection and connection capability. See "<u>Automatic Discovery and Connection to</u> <u>Targets</u>" section.

Automatic Discovery and Connection to Targets

- Make sure the ib_srp module is loaded, the SRP Initiator can reach an SRP Target, and that an SM is running.
- To connect to all the existing Targets in the fabric, run "srp_daemon -e -o". This utility will scan the fabric once, connect to every Target it detects, and then exit.



srp_daemon will follow the configuration it finds in
/etc/srp_daemon.conf. Thus, it will ignore a target that is
disallowed in the configuration file.

- To connect to all the existing Targets in the fabric and to connect to new targets that will join the fabric, execute srp_daemon -e. This utility continues to execute until it is either killed by the user or encounters connection errors (such as no SM in the fabric).
- To execute SRP daemon as a daemon on all the ports:
 - srp_daemon.sh (found under /usr/sbin/). srp_daemon.sh sends its log to /var/log/srp_daemon.log.

• Start the srpd service script, run service srpd start

For the changes in openib.conf to take effect, run:

/etc/init.d/openibd restart

Multiple Connections from Initiator InfiniBand Port to the Target

Some system configurations may need multiple SRP connections from the SRP Initiator to the same SRP Target: to the same Target IB port, or to different IB ports on the same Target HCA.

In case of a single Target IB port, i.e., SRP connections use the same path, the configuration is enabled using a different initiator_ext value for each SRP connection. The initiator_ext value is a 16-hexadecimal-digit value specified in the connection command.

Also in case of two physical connections (i.e., network paths) from a single initiator IB port to two different IB ports on the same Target HCA, there is need for a different initiator_ext value on each path. The conventions is to use the Target port GUID as the initiator_ext value for the relevant path.

If you use srp_daemon with -n flag, it automatically assigns initiator_ext values according to this convention. For example:

id_ext=200500A0B81146A1,ioc_guid=0002c90200402bec,\
dgid=fe800000000000000002c90200402bed,pkey=ffff,\
service_id=200500a0b81146a1,initiator_ext=ed2b400002c90200

Notes:

- It is recommended to use the -n flag for all srp_daemon invocations.
- ibsrpdm does not have a corresponding option.

• srp_daemon.sh always uses the -n option (whether invoked manually by the user, or automatically at startup by setting SRP_DAEMON_ENABLE to yes).

High Availability (HA)

High Availability works using the Device-Mapper (DM) multipath and the SRP daemon. Each initiator is connected to the same target from several ports/HCAs. The DM multipath is responsible for joining together different paths to the same target and for failover between paths when one of them goes offline. Multipath will be executed on newly joined SCSI devices.

Each initiator should execute several instances of the SRP daemon, one for each port. At startup, each SRP daemon detects the SRP Targets in the fabric and sends requests to the ib_srp module to connect to each of them. These SRP daemons also detect targets that subsequently join the fabric, and send the ib_srp module requests to connect to them as well.

Operation

When a path (from port1) to a target fails, the ib_srp module starts an error recovery process. If this process gets to the reset_host stage and there is no path to the target from this port, ib_srp will remove this scsi_host. After the scsi_host is removed, multipath switches to another path to this target (from another port/HCA).

When the failed path recovers, it will be detected by the SRP daemon. The SRP daemon will then request ib_srp to connect to this target. Once the connection is up, there will be a new scsi_host for this target. Multipath will be executed on the devices of this host, returning to the original state (prior to the failed path).

Manual Activation of High Availability

Initialization - execute after each boot of the driver:

- 1. Execute modprobe dm-multipath
- 2. Execute modprobe ib-srp
- 3. Make sure you have created file /etc/udev/rules.d/91-srp.rules as described above
- 4. Execute for each port and each HCA:

```
srp_daemon -c -e -R 300 -i <InfiniBand HCA name> -p <port
number>
```

This step can be performed by executing srp_daemon.sh, which sends its log to /var/log/srp_daemon.log.

Now it is possible to access the SRP LUNs on /dev/mapper/.

(i) Note

It is possible for regular (non-SRP) LUNs to also be present; the SRP LUNs may be identified by their names. You can configure the /etc/multipath.conf file to change multipath behavior.

) Note

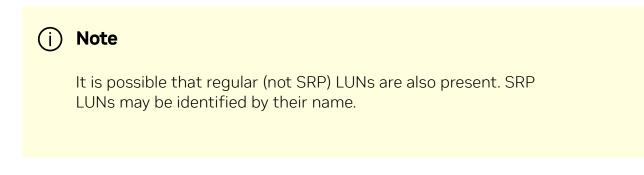
It is also possible that the SRP LUNs will not appear under /dev/mapper/. This can occur if the SRP LUNs are in the black-list of multipath. Edit the 'blacklist' section in /etc/multipath.conf and make sure the SRP LUNs are not blacklisted.

Automatic Activation of High Availability

• Start srpd service, run:

service srpd start

• From the next loading of the driver it will be possible to access the SRP LUNs on /dev/mapper/



• It is possible to see the output of the SRP daemon in /var/log/srp_daemon.log

Shutting Down SRP

SRP can be shutdown by using "rmmod ib_srp", or by stopping the OFED driver ("/etc/init.d/openibd stop"), or as a by-product of a complete system shutdown.

Prior to shutting down SRP, remove all references to it. The actions you need to take depend on the way SRP was loaded. There are three cases:

1. Without High Availability

When working without High Availability, you should unmount the SRP partitions that were mounted prior to shutting down SRP.

2. After Manual Activation of High Availability

If you manually activated SRP High Availability, perform the following steps:

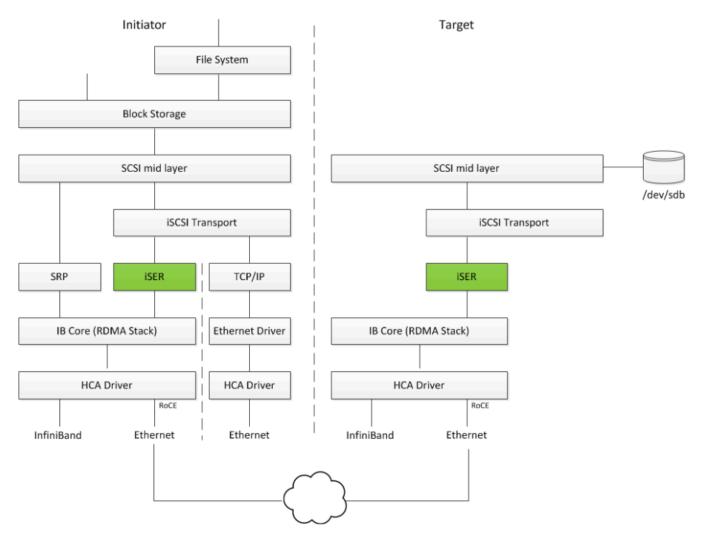
- 1. Unmount all SRP partitions that were mounted.
- 2. Stop service srpd (Kill the SRP daemon instances).
- 3. Make sure there are no multipath instances running. If there are multiple instances, wait for them to end or kill them.
- 4. Run: multipath -F
- 3. After Automatic Activation of High Availability

If SRP High Availability was automatically activated, SRP shutdown must be part of the driver shutdown ("/etc/init.d/openibd stop") which performs Steps 2-4 of case b above. However, you still have to unmount all SRP partitions that were mounted before driver shutdown.

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.

iSER uses the RDMA protocol suite to supply higher bandwidth for block storage transfers (zero time copy behavior). To that fact, it eliminates the TCP/IP processing overhead while preserving the compatibility with iSCSI protocol.



There are three target implementation of ISER:

• Linux SCSI target framework (tgt)

- Linux-IO target (LIO)
- Generic SCSI target subsystem for Linux (SCST)

Each one of those targets can work in TCP or iSER transport modes.

iSER also supports RoCE without any additional configuration required. To bond the RoCE interfaces, set the fail_over_mac option in the bonding driver (see "<u>Bonding IPoIB</u>").

RDMA/RoCE is located below the iSER block on the network stack. In order to run iSER, the RDMA layer should be configured and validated (over Ethernet or InfiniBand). For troubleshooting RDMA, please refer to "<u>HowTo Enable, Verify and Troubleshoot RDMA</u>" on the Community website.

iSER Initiator

The iSER initiator is controlled through the iSCSI interface available from the iscsiinitiator-utils package.

To discover and log into iSCSI targets, as well as access and manage the open-iscsi database use the iscasiadm utility, a command-line tool.

To enable iSER as a transport protocol use "I iser" as a parameter of the iscasiadm command.

Example for discovering and connecting targets over iSER:

```
iscsiadm -m discovery -o new -o old -t st -I iser -p <ip:port> -l
```

Note that the target implementation (e.g. LIO, SCST, TGT) does not affect he initiator process and configuration.

iSER Targets

j) Note

Setting the iSER target is out of scope of this manual. For guidelines of how to do so, please refer to the relevant target documentation (e.g. stgt, targetcli).

Targets settings such as timeouts and retries are set the same as any other iSCSI targets.



If targets are set to auto connect on boot, and targets are unreachable, it may take a long time to continue the boot process if timeouts and max retries are set too high.

For various configuration, troubleshooting and debugging examples, refer to <u>Storage</u> <u>Solutions</u> on the Community website.

Lustre

Lustre is an open-source, parallel distributed file system, generally used for large-scale cluster computing that supports many requirements of leadership class HPC simulation environments.

Lustre Compilation for MLNX_OFED:



This procedure applies to RHEL/SLES OSs supported by Lustre. For further information, please refer to Lustre Release Notes.

To compile Lustre version 2.4.0 and higher:

\$./configure --with-o2ib=/usr/src/ofa_kernel/default/
\$ make rpms

\$ EXTRA_LNET_INCLUDE="-I/usr/src/ofa_kernel/default/include/-include /usr/src/ofa_kernel/default/include/linux/compat-2.6.h" ./configure --witho2ib=/usr/src/ofa_kernel/default/ \$ EXTRA_LNET_INCLUDE="-I/usr/src/ofa_kernel/default/include/-include /usr/src/ofa_kernel/default/include/linux/compat-2.6.h" make rpms

For full installation example, refer to <u>HowTo Install NVIDIA OFED driver for Lustre</u> Community post.

NVME-oF - NVM Express over Fabrics

NVME-oF

NVME-oF enables NVMe message-based commands to transfer data between a host computer and a target solid-state storage device or system over a network such as Ethernet, Fibre Channel, and InfiniBand. Tunneling NVMe commands through an RDMA fabric provides a high throughput and a low latency.

For information on how to configure NVME-oF, please refer to the <u>HowTo Configure NVMe</u> <u>over Fabrics</u> Community post.



The --with-nvmf installation option should **not** be specified, if nvmetcp kernel module is used. In this case, the native Inbox nvme-tcp kernel module will be loaded.

NVME-oF Target Offload

(i) Note

This feature is only supported for ConnectX-5 adapter cards family and above.

NVME-oF Target Offload is an implementation of the new NVME-oF standard Target (server) side in hardware. Starting from ConnectX-5 family cards, all regular IO requests can be processed by the HCA, with the HCA sending IO requests directly to a real NVMe PCI device, using peer-to-peer PCI communications. This means that excluding connection management and error flows, no CPU utilization will be observed during NVME-oF traffic.

- For instructions on how to configure NVME-oF target offload, refer to <u>HowTo</u> <u>Configure NVME-oF Target Offload</u> Community post.
- For instructions on how to verify that NVME-oF target offload is working properly, refer to <u>Simple NVMe-oF Target Offload Benchmark</u> Community post.

Virtualization

The chapter contains the following sections:

- Single Root IO Virtualization (SR-IOV)
- Enabling Paravirtualization
- VXLAN Hardware Stateless Offloads
- Q-in-Q Encapsulation per VF in Linux (VST)
- 802.1Q Double-Tagging
- <u>Scalable Functions</u>

Single Root IO Virtualization (SR-IOV)

Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus. This technology enables multiple virtual instances of the device with separate resources. NVIDIA adapters are capable of exposing up to 127 virtual instances (Virtual Functions (VFs) for each port in the NVIDIA ConnectX® family cards. These virtual functions can then be provisioned separately. Each VF can be seen as an additional device connected to the Physical Function. It shares the same resources with the Physical Function, and its number of ports equals those of the Physical Function.

SR-IOV is commonly used in conjunction with an SR-IOV enabled hypervisor to provide virtual machines direct hardware access to network resources hence increasing its performance.

In this chapter we will demonstrate setup and configuration of SR-IOV in a Red Hat Linux environment using ConnectX[®] VPI adapter cards.

System Requirements

To set up an SR-IOV environment, the following is required:

- MLNX_OFED Driver
- A server/blade with an SR-IOV-capable motherboard BIOS
- Hypervisor that supports SR-IOV such as: Red Hat Enterprise Linux Server Version 6
- NVIDIA ConnectX[®] VPI Adapter Card family with SR-IOV capability

Setting Up SR-IOV

Depending on your system, perform the steps below to set up your BIOS. The figures used in this section are for illustration purposes only. For further information, please refer to the appropriate BIOS User Manual:

1. Enable "SR-IOV" in the system BIOS.

Advanced PCI/PnP Settings			Options Disabled Enabled	
WARNING: Setting wrong values in below sections may cause system to malfunction.				
	EDisabled] EPXE1	↔ †1 +- F1 F10 ESC		

2. Enable "Intel Virtualization Technology".

		IOS SETUP UTILITY				
Advanced						
Microcode Rev	:14			When	enabled, a UMM	
Cache L1	:256 KB			can i	utilize the	
Cache L2	:1024 KB			addit	tional HW Caps	
Cache L3	:12288 KB			prov	ided by Intel(R)	
Ratio Status	:Unlocked (M	in:12, Max:18)		Virt	alization Tech.	
Ratio Actual Value:18				Note: A full reset is		
				requi	ired to change	
CPII Ratin		[Auto]		the s	setting.	
C1E Support		[Enabled]				
Hardware Prefetcher [Enabled]						
Adjacent Cache Line Prefetch [Enabled]						
DCII Prefetcher		[Enabled]				
Data Reuse Optimization		[Enabled]		↔	Select Screen	
		[Modern ordering]		† ‡	Select Item	
Intel(R) Virtual		[Enabled]		4-	Change Option	
Execute-Disable Bit Capability				F1		
Intel AES-NI		IDisabledl		F10	Save and Exit	
Simultaneous Mul				ESC	Exit	
Active Processon		IAIII				
Intel(R) EIST Te	echno logy	lEnabledI	•			
	C) Comunicable	1985-2009, American	Man			

- 3. Install a hypervisor that supports SR-IOV.
- 4. Depending on your system, update the /boot/grub/grub.conf file to include a similar command line load parameter for the Linux kernel.

For example, to Intel systems, add:

```
default=0
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
```

```
hiddenmenu
title Red Hat Enterprise Linux Server (4.x.x)
    root (hd0,0)
    kernel /vmlinuz-4.x.x ro
root=/dev/VolGroup00/LogVol00 rhgb quiet
    intel_iommu=on initrd /initrd-4.x.x.img
```

Note: Please make sure the parameter "intel_iommu=on" exists when updating the /boot/grub/grub.conf file, otherwise SR-IOV cannot be loaded.

Some OSs use /boot/grub2/grub.cfg file. If your server uses such file, please edit this file instead (add "intel_iommu=on" for the relevant menu entry at the end of the line that starts with "linux16").

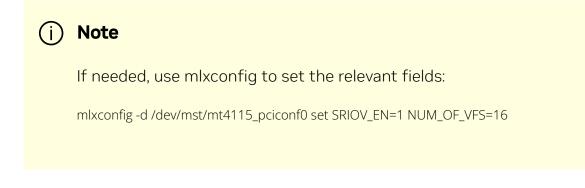
Configuring SR-IOV (Ethernet)

To set SR-IOV in Ethernet mode, refer to <u>HowTo Configure SR-IOV for ConnectX-</u> <u>4/ConnectX- 5/ConnectX-6 with KVM (Ethernet)</u> Community Post.

Configuring SR-IOV (InfiniBand)

- 1. Install the MLNX_OFED driver for Linux that supports SR-IOV.
- 2. Check if SR-IOV is enabled in the firmware.

```
mlxconfig -d /dev/mst/mt4115_pciconf0 q
Device #1:
Device type: Connect4
PCI device: /dev/mst/mt4115_pciconf0
Configurations: Current
SRIOV_EN 1
NUM_OF_VFS 8
```



- 3. Reboot the server.
- 4. Write to the sysfs file the number of Virtual Functions you need to create for the PF. You can use one of the following equivalent files:

You can use one of the following equivalent files:

- A standard Linux kernel generated file that is available in the new kernels.

```
echo [num_vfs] >
/sys/class/infiniband/mlx5_0/device/sriov_numvfs
```

Note: This file will be generated only if IOMMU is set in the grub.conf file (by adding intel_iommu=on, as seen in the fourth step under "<u>Setting Up SR-IOV</u>").

- A file generated by the mlx5_core driver with the same functionality as the kernel generated one.

```
echo [num_vfs] >
/sys/class/infiniband/mlx5_0/device/mlx5_num_vfs
```

Note: This file is used by old kernels that do not support the standard file. In such kernels, using sriov_numvfs results in the following error: "bash: echo: write error: Function not implemented".

The following rules apply when writing to these files:

- If there are no VFs assigned, the number of VFs can be changed to any valid value (0 - max #VFs as set during FW burning)

- If there are VFs assigned to a VM, it is not possible to change the number of VFs

- If the administrator unloads the driver on the PF while there are no VFs assigned, the driver will unload and SRI-OV will be disabled

- If there are VFs assigned while the driver of the PF is unloaded, SR-IOV will not be disabled. This means that VFs will be visible on the VM. However, they will not be operational. This is applicable to OSs with kernels that use pci_stub and not vfio.

- The VF driver will discover this situation and will close its resources

- When the driver on the PF is reloaded, the VF becomes operational. The administrator of the VF will need to restart the driver in order to resume working with the VF.

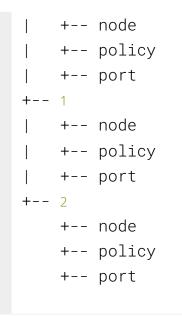
5. Load the driver. To verify that the VFs were created. Run:

```
lspci | grep Mellanox
08:00.0 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4]
08:00.1 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4]
08:00.2 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4 Virtual Function]
08:00.3 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4 Virtual Function]
08:00.4 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4 Virtual Function]
08:00.5 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4 Virtual Function]
08:00.5 Infiniband controller: Mellanox Technologies MT27700
Family [ConnectX-4 Virtual Function]
```

6. Configure the VFs.

After VFs are created, 3 sysfs entries per VF are available under /sys/class/infiniband/mlx5_<PF INDEX>/device/sriov (shown below for VFs 0 to 2):

+-- 0



For each Virtual Function, the following files are available:

- Node - Node's GUID:

The user can set the node GUID by writing to the /sys/class/infiniband/<PF>/device/sriov/<index>/node file. The example below, shows how to set the node GUID for VF 0 of mlx5_0.

```
echo 00:11:22:33:44:55:1:0 >
/sys/class/infiniband/mlx5_0/device/sriov/0/node
```

- Port - Port's GUID:

The user can set the port GUID by writing to the /sys/class/infiniband/<PF>/device/sriov/<index>/port file. The example below, shows how to set the port GUID for VF 0 of mlx5_0.

```
echo 00:11:22:33:44:55:2:0 >
/sys/class/infiniband/mlx5_0/device/sriov/0/port
```

- Policy - The vport's policy. The user can set the port GUID by writing to the /sys/class/infiniband/<PF>/device/sriov/<index>/port file. The policy can be one of:

- Down - the VPort PortState remains 'Down'

- Up - if the current VPort PortState is 'Down', it is modified to 'Initialize'. In all other states, it is unmodified. The result is that the SM may bring the VPort up.

- Follow - follows the PortState of the physical port. If the PortState of the physical port is 'Active', then the VPort implements the 'Up' policy. Otherwise, the VPort PortState is 'Down'.

Notes:

- The policy of all the vports is initialized to "Down" after the PF driver is restarted except for VPort0 for which the policy is modified to 'Follow' by the PF driver.

- To see the VFs configuration, you must unbind and bind them or reboot the VMs if the VFs were assigned.

7. Make sure that OpenSM supports Virtualization (Virtualization must be enabled).

The /etc/opensm/opensm.conf file should contain the following line:

virt_enabled 2

Note: OpenSM and any other utility that uses SMP MADs (ibnetdiscover, sminfo, iblink- info, smpdump, ibqueryerr, ibdiagnet and smpquery) should run on the PF and not on the VFs. In case of multi PFs (multi-host), OpenSM should run on HostO.

VFs Initialization Note

Since the same mlx5_core driver supports both Physical and Virtual Functions, once the Virtual Functions are created, the driver of the PF will attempt to initialize them so they will be available to the OS owning the PF. If you want to assign a Virtual Function to a VM, you need to make sure the VF is not used by the PF driver. If a VF is used, you should first unbind it before assigning to a VM.

To unbind a device use the following command:

1. Get the full PCI address of the device.

lspci -D

Example:

0000:09:00.2

2. Unbind the device.

echo 0000:09:00.2 > /sys/bus/pci/drivers/mlx5_core/unbind

3. Bind the unbound VF.

echo 0000:09:00.2 > /sys/bus/pci/drivers/mlx5_core/bind

PCI BDF Mapping of PFs and VFs

PCI addresses are sequential for both of the PF and their VFs. Assuming the card's PCI slot is 05:00 and it has 2 ports, the PFs PCI address will be 05:00.0 and 05:00.1.

Given 3 VFs per PF, the VFs PCI addresses will be:

```
05:00.2-4 for VFs 0-2 of PF 0 (mlx5_0)
05:00.5-7 for VFs 0-2 of PF 1 (mlx5_1)
```

Additional SR-IOV Configurations

Assigning a Virtual Function to a Virtual Machine

This section describes a mechanism for adding a SR-IOV VF to a Virtual Machine.

Assigning the SR-IOV Virtual Function to the Red Hat KVM VM Server

- 1. Run the virt-manager.
- 2. Double click on the virtual machine and open its Properties.
- 3. Go to Details \rightarrow Add hardware \rightarrow PCI host device.



- 4. Choose a NVIDIA virtual function according to its PCI device (e.g., 00:03.1)
- 5. If the Virtual Machine is up reboot it, otherwise start it.
- 6. Log into the virtual machine and verify that it recognizes the NVIDIA card. Run:

lspci | grep Mellanox

Example:

```
lspci | grep Mellanox
01:00.0 Infiniband controller: Mellanox Technologies MT28800
Family [ConnectX-5 Ex]
```

7. Add the device to the /etc/sysconfig/network-scripts/ifcfg-ethX configuration file. The MAC address for every virtual function is configured randomly, therefore it is not necessary to add it.

Ethernet Virtual Function Configuration when Running SR-IOV

SR-IOV Virtual function configuration can be done through Hypervisor iprout2/netlink tool, if present. Otherwise, it can be done via sysfs.

```
ip link set { dev DEVICE | group DEVGROUP } [ { up | down } ]
...
[ vf NUM [ mac LLADDR ] [ vlan VLANID [ qos VLAN-QOS ] ]
...
[ spoofchk { on | off} ] ]
...
sysfs configuration (ConnectX-4):
/sys/class/net/enp8s0f0/device/sriov/[VF]
+-- [VF]
| +-- config
```



VLAN Guest Tagging (VGT) and VLAN Switch Tagging (VST)

When running ETH ports on VGT, the ports may be configured to simply pass through packets as is from VFs (VLAN Guest Tagging), or the administrator may configure the Hypervisor to silently force packets to be associated with a VLAN/Qos (VLAN Switch Tagging).

In the latter case, untagged or priority-tagged outgoing packets from the guest will have the VLAN tag inserted, and incoming packets will have the VLAN tag removed.

The default behavior is VGT.

To configure VF VST mode, run:

ip link set dev <PF device> vf <NUM> vlan <vlan_id> [qos <qos>]

where:

- NUM = 0..max-vf-num
- vlan_id = 0..4095
- qos = 0..7

For example:

- ip link set dev eth2 vf 2 vlan 10 qos 3 sets VST mode for VF #2 belonging to PF eth2, with vlan_id = 10 and qos = 3
- ip link set dev eth2 vf 2 vlan 0 sets mode for VF 2 back to VGT

Additional Ethernet VF Configuration Options

 Guest MAC configuration - by default, guest MAC addresses are configured to be all zeroes. If the administrator wishes the guest to always start up with the same MAC, he/she should configure guest MACs before the guest driver comes up. The guest MAC may be configured by using:

ip link set dev <PF device> vf <NUM> mac <LLADDR>

For legacy and ConnectX-4 guests, which do not generate random MACs, the administrator should always configure their MAC addresses via IP link, as above.

• **Spoof checking** - Spoof checking is currently available only on upstream kernels newer than 3.1.

ip link set dev <PF device> vf <NUM> spoofchk [on | off]

• Guest Link State

ip link set dev <PF device> vf <UM> state [enable| disable| auto]

Virtual Function Statistics

Virtual function statistics can be queried via sysfs:

```
cat /sys/class/infiniband/mlx5_2/device/sriov/2/stats tx_packets :
5011
tx_bytes : 4450870
tx_dropped : 0
rx_packets : 5003
rx_bytes : 4450222
rx_broadcast : 0
rx_multicast : 0
tx_broadcast : 0
tx_multicast : 8
rx_dropped : 0
```

Mapping VFs to Ports

To view the VFs mapping to ports:

Use the ip link tool v2.6.34~3 and above.

```
ip link
```

Output:

```
61: p1p1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN
mode DEFAULT group default qlen 1000
    link/ether 00:02:c9:f1:72:e0 brd ff:ff:ff:ff:ff
    vf 0 MAC 00:00:00:00:00;00, vlan 4095, spoof checking off, link-
state auto
    vf 37 MAC 00:00:00:00:00:00, vlan 4095, spoof checking off, link-
state auto
    vf 38 MAC ff:ff:ff:ff:ff:ff; vlan 65535, spoof checking off,
link-state disable
```

```
vf 39 MAC ff:ff:ff:ff:ff, vlan 65535, spoof checking off, link-state disable
```

However, even VFs that are not assigned to the net device, could be used to set and change its settings. For example, the following is a valid command to change the spoof check:

```
ip link set dev p1p1 vf 38 spoofchk on
```

This command will affect only the **vf38**. The changes can be seen in ip link on the net device that this device is assigned to.

RoCE Support

RoCE is supported on Virtual Functions and VLANs may be used with it. For RoCE, the hypervisor GID table size is of 16 entries while the VFs share the remaining 112 entries. When the number of VFs is larger than 56 entries, some of them will have GID table with only a single entry which is inadequate if VF's Ethernet device is assigned with an IP address.

Virtual Guest Tagging (VGT+)

VGT+ is an advanced mode of Virtual Guest Tagging (VGT), in which a VF is allowed to tag its own packets as in VGT, but is still subject to an administrative VLAN trunk policy. The policy determines which VLAN IDs are allowed to be transmitted or received. The policy does not determine the user priority, which is left unchanged.

Packets can be sent in one of the following modes: when the VF is allowed to send/receive untagged and priority tagged traffic and when it is not. No default VLAN is defined for

VGT+ port. The send packets are passed to the eSwitch only if they match the set, and the received packets are forwarded to the VF only if they match the set.

Configuration

(i) Note

When working in SR-IOV, the default operating mode is VGT.

To enable VGT+ mode:

Set the corresponding port/VF (in the example below port eth5, VFO) range of allowed VLANs.

```
echo "<add> <start_vid> <end_vid>" > /sys/class/net/eth5/device/sriov/0/trunk
```

Examples:

• Adding VLAN ID range (4-15) to trunk:

echo add 4 15 > /sys/class/net/eth5/device/sriov/0/trunk

• Adding a single VLAN ID to trunk:

echo add 17 17 > /sys/class/net/eth5/device/sriov/0/trunk

Note: When VLAN ID = 0, it indicates that untagged and priority-tagged traffics are allowed

To disable VGT+ mode, make sure to remove all VLANs.

echo rem 0 4095 > /sys/class/net/eth5/device/sriov/0/trunk

To remove selected VLANs.

• Remove VLAN ID range (4-15) from trunk:

echo rem 4 15 > /sys/class/net/eth5/device/sriov/0/trunk

• Remove a single VLAN ID from trunk:

```
echo rem 17 17 > /sys/class/net/eth5/device/sriov/0/trunk
```

SR-IOV Advanced Security Features

SR-IOV MAC Anti-Spoofing

Normally, MAC addresses are unique identifiers assigned to network interfaces, and they are fixed addresses that cannot be changed. MAC address spoofing is a technique for altering the MAC address to serve different purposes. Some of the cases in which a MAC address is altered can be legal, while others can be illegal and abuse security mechanisms or disguises a possible attacker.

The SR-IOV MAC address anti-spoofing feature, also known as MAC Spoof Check provides protection against malicious VM MAC address forging. If the network administrator assigns a MAC address to a VF (through the hypervisor) and enables spoof check on it, this will limit the end user to send traffic only from the assigned MAC address of that VF.

MAC Anti-Spoofing Configuration



MAC anti-spoofing is disabled by default.

In the configuration example below, the VM is located on VF-0 and has the following MAC address: 11:22:33:44:55:66.

There are two ways to enable or disable MAC anti-spoofing:

- 1. Use the standard IP link commands available from Kernel 3.10 and above.
 - 1. To enable MAC anti-spoofing, run:

ip link set ens785f1 vf 0 spoofchk on

2. To disable MAC anti-spoofing, run:

ip link set ens785f1 vf 0 spoofchk off

- 2. Specify echo "ON" or "OFF" to the file located under /sys/class/net/<ifname / device/sriov/<VF index>/spoofcheck.
 - 1. To enable MAC anti-spoofing, run:

echo "ON" > /sys/class/net/ens785f1/vf/0/spoofchk

2. To disable MAC anti-spoofing, run:

echo "OFF" > /sys/class/net/ens785f1/vf/0/spoofchk

(i) Note

This configuration is non-persistent and does not survive driver restart.

Limit and Bandwidth Share Per VF

This feature enables rate limiting traffic per VF in SR-IOV mode. For details on how to configure rate limit per VF for ConnectX-4 and above adapter cards, please refer to <u>HowTo</u> <u>Configure Rate Limit per VF for ConnectX-4/ConnectX-5/ConnectX-6</u> Community post.

Limit Bandwidth per Group of VFs

VFs Rate Limit for vSwitch (OVS) feature allows users to join available VFs into groups and set a rate limitation on each group. Rate limitation on a VF group ensures that the total Tx bandwidth that the VFs in this group get (altogether combined) will not exceed the given value.

With this feature, a VF can still be configured with an individual rate limit as in the past (under /sys/class/net//device/sriov//max_tx_rate). However, the actual bandwidth limit on the VF will eventually be determined considering the VF group limitation and how many VFs are in the same group.

For example: 2 VFs (0 and 1) are attached to group 3.

Case 1: The rate limitation on the group is set to 20G. Rate limit of each VF is 15G

Result: Each VF will have a rate limit of 10G

Case 2: Group's max rate limitation is still set to 20G. VF 0 is configured to 30G limit, while VF 1 is configured to 5G rate limit

Result: VF 0 will have 15G de-facto. VF 1 will have 5G

The rule of thumb is that the group's bandwidth is distributed evenly between the number of VFs in the group. If there are leftovers, they will be assigned to VFs whose individual

rate limit has not been met yet.

VFs Rate Limit Feature Configuration

- 1. When VF rate group is supported by FW, the driver will create a new hierarchy in the SRI-OV sysfs named "groups" (/sys/class/net/<ifname>/device/sriov/groups/). It will contain all the info and the configurations allowed for VF groups.
- 2. All VFs are placed in group 0 by default since it is the only existing group following the initial driver start. It would be the only group available under /sys/class/net/<ifname>/device/sriov/groups/
- 3. The VF can be moved to a different group by writing to the group file -> echo \$GROUP_ID > /sys/class/net/<ifname>/device/sriov/<vf_id>/group
- 4. The group IDs allowed are 0-255
- 5. Only when there is at least 1 VF in a group, there will be a group configuration available under /sys/class/net/<ifname>/device/sriov/groups/ (Except for group 0, which is always available even when it's empty).
- 6. Once the group is created (by moving at least 1 VF to that group), users can configure the group's rate limit. For example:
 - 1. echo 10000 > /sys/class/net/<ifname>/device/sriov/5/max_tx_rate setting individual rate limitation of VF 5 to 10G (Optional)
 - 2. echo 7 > /sys/class/net/<ifname>/device/sriov/5/group moving VF 5 to group 7
 - 3. echo 5000 > /sys/class/net/<ifname>/device/sriov/groups/7/max_tx_rate setting group 7 with rate limitation of 5G
 - 4. When running traffic via VF 5 now, it will be limited to 5G because of the group rate limit even though the VF itself is limited to 10G
 - 5. echo 3 > /sys/class/net/<ifname>/device/sriov/5/group moving VF 5 to group 3
 - 6. Group 7 will now disappear from /sys/class/net/<ifname>/device/sriov/groups since there are 0 VFs in it. Group 3 will now appear. Since there's no rate limit on group 3, VF 5 can transmit at 10G (thanks to its individual configuration)

Notes:

- You can see to which group the VF belongs to in the 'stats' sysfs (cat /sys/class/net/<ifname>/device/sriov/<vf_num>/stats)
- You can see the current rate limit and number of attached VFs to a group in the group's 'config' sysfs (cat /sys/class/net/<ifname>/device/sriov/groups/<group_id>/config)

Bandwidth Guarantee per Group of VFs

Bandwidth guarantee (minimum BW) can be set on a group of VFs to ensure this group is able to transmit <u>at least</u> the amount of bandwidth specified on the wire.

Note the following:

- The minimum BW settings on VF groups determine how the groups share the total BW between themselves. It does not impact an individual VF's rate settings.
- The total minimum BW that is set on the VF groups should not exceed the total line rate. Otherwise, results are unexpected.
- It is still possible to set minimum BW on the individual VFs inside the group. This will determine how the VFs share the group's minimum BW between themselves. The total minimum BW of the VF member should not exceed the minimum BW of the group.

For instruction on how to create groups of VFs, see <u>Limit Bandwidth per Group of VFs</u> above.

Example

With a 40Gb link speed, assuming 4 groups and default group 0 have been created:

```
echo 20000 >
/sys/class/net/<ifname>/device/sriov/group/1/min_tx_rate
echo 5000 > /sys/class/net/<ifname>/device/sriov/group/2/min_tx_rate
echo 15000 >
/sys/class/net/<ifname>/device/sriov/group/3/min_tx_rate
```

```
Group 0(default) : 0 - No BW guarantee is configured.
Group 1 : 20000 - This is the maximum min rate among groups
Group 2 : 5000 which is 25% of the maximum min rate
Group 3 : 15000 which is 75% of the maximum min rate
Group 4 : 0 - No BW guarantee is configured.
```

Assuming there are VFs attempting to transmit in full line rate in all groups, the results would look like: In which case, the minimum BW allocation would be:

Group0 - Will have no BW to use since no BW guarantee was set on it while other groups do have such settings. Group1 - Will transmit at 20Gb/s Group2 - Will transmit at 5Gb/s Group3 - Will transmit at 15Gb/s Group4 - Will have no BW to use since no BW guarantee was set on it while other groups do have such settings.

Privileged VFs

In case a malicious driver is running over one of the VFs, and in case that VF's permissions are not restricted, this may open security holes. However, VFs can be marked as trusted and can thus receive an exclusive subset of physical function privileges or permissions. For example, in case of allowing all VFs, rather than specific VFs, to enter a promiscuous mode as a privilege, this will enable malicious users to sniff and monitor the entire physical port for incoming traffic, including traffic targeting other VFs, which is considered a severe security hole.

Privileged VFs Configuration

In the configuration example below, the VM is located on VF-0 and has the following MAC address: 11:22:33:44:55:66.

There are two ways to enable or disable trust:

- 1. Use the standard IP link commands available from Kernel 4.5 and above.
 - 1. To enable trust for a specific VF, run:

ip link set ens785f1 vf 0 trust on

2. To disable trust for a specific VF, run:

ip link set ens785f1 vf 0 trust off

- 2. Specify echo "ON" or "OFF" to the file located under /sys/class/net/<ETH_IF_NAME> / device/sriov/<VF index>/trust.
- 1. To enable trust for a specific VF, run:

```
echo "ON" > /sys/class/net/ens785f1/device/sriov/0/trust
```

2. To disable trust for a specific VF, run:

echo "OFF" > /sys/class/net/ens785f1/device/sriov/0/trust

Probed VFs

Probing Virtual Functions (VFs) after SR-IOV is enabled might consume the adapter cards' resources. Therefore, it is recommended not to enable probing of VFs when no monitoring of the VM is needed.

VF probing can be disabled in two ways, depending on the kernel version installed on your server:

- 1. If the kernel version installed is v4.12 or above, it is recommended to use the PCI sysfs interface sriov_drivers_autoprobe. For more information, see <u>linux-next branch</u>.
- 2. If the kernel version installed is older than v4.12, it is recommended to use the mlx5_core module parameter probe_vf with driver version 4.1 or above.

Example:

echo 0 > /sys/module/mlx5_core/parameters/probe_vf

For more information on how to probe VFs, see <u>HowTo Configure and Probe VFs on mlx5</u> <u>Drivers</u>Community post.

VF Promiscuous Rx Modes

VF Promiscuous Mode

VFs can enter a promiscuous mode that enables receiving the unmatched traffic and all the multicast traffic that reaches the physical port in addition to the traffic originally targeted to the VF. The unmatched traffic is any traffic's DMAC that does not match any of the VFs' or PFs' MAC addresses.

Note: Only privileged/trusted VFs can enter the VF promiscuous mode.

To set the promiscuous mode on for a VF, run:

ifconfig eth2 promisc

To exit the promiscuous mode, run:

≻

ifconfig eth2 -promisc

VF All-Multi Mode

VFs can enter an all-multi mode that enables receiving all the multicast traffic sent from/to the other functions on the same physical port in addition to the traffic originally targeted to the VF.

Note: Only privileged/trusted VFs can enter the all-multi RX mode.

To set the all-multi mode on for a VF, run:

ifconfig eth2 allmulti

To exit the all-multi mode, run:

≽

≻

#ifconfig eth2 -allmulti

Uninstalling the SR-IOV Driver

To uninstall SR-IOV driver, perform the following:

≻

1. For Hypervisors, detach all the Virtual Functions (VF) from all the Virtual Machines (VM) or stop the Virtual Machines that use the Virtual Functions.

Please be aware that stopping the driver when there are VMs that use the VFs, will cause machine to hang.

2. Run the script below. Please be aware, uninstalling the driver deletes the entire driver's file, but does not unload the driver.

[root@swl022 ~]# /usr/sbin/ofed_uninstall.sh This program will uninstall all OFED packages on your machine. Do you want to continue?[y/N]:y Running /usr/sbin/vendor_pre_uninstall.sh Removing OFED Software installations Running /bin/rpm -e --allmatches kernel-ib kernel-ib-devel libibverbs libibverbs-devel libibverbs-devel-static libibverbsutils libmlx4 libmlx4-devel libibcm libibcm-devel libibumad libibumad-devel libibumad-static libibmad libibmad-devel libibmad-static librdmacm librdmacm-utils librdmacm-devel ibacm opensm-libs opensm-devel perftest compat-dapl compat-dapldevel dapl dapl-devel dapl-devel-static dapl-utils srptools infiniband-diags-quest ofed-scripts opensm-devel warning: /etc/infiniband/openib.conf saved as /etc/infiniband/openib.conf.rpmsave Running /tmp/2818-ofed_vendor_post_uninstall.sh

3. Restart the server.

SR-IOV Live Migration

j) Note

This feature is supported in Ethernet mode only.

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 other adverse effects for the end user.

Using the Migration process is useful for:

- load balancing
- hardware independence
- energy saving
- geographic migration
- fault tolerance

Migration works by sending the state of the guest virtual machine's memory and any virtualized devices to a destination host physical machine. Migrations can be performed live or not, in the live case, the migration will not disrupt the user operations and it will be transparent to it as explained in the sections below.

Non-Live Migration

When using the non-live migration process, the Hypervisor suspends the guest virtual machine, then moves an image of the guest virtual machine's memory to the destination host physical machine. The guest virtual machine is then resumed on the destination host physical machine, and the memory the guest virtual machine used on the source host physical machine is freed. The time it takes to complete such a migration depends on the network bandwidth and latency. If the network is experiencing heavy use or low bandwidth, the migration will take longer then desired.

Live Migration

When using the Live Migration process, the guest virtual machine continues to run on the source host physical machine while its memory pages are transferred to the destination host physical machine. During migration, the Hypervisor monitors the source for any changes in the pages it has already transferred and begins to transfer these changes when all of the initial pages have been transferred.

It also estimates transfer speed during migration, so when the remaining amount of data to transfer will take a certain configurable period of time, it will suspend the original guest virtual machine, transfer the remaining data, and resume the same guest virtual machine on the destination host physical machine.

MLX5 VF Live Migration

The purpose of this section is to demonstrate how to perform basic live migration of a QEMU VM with an MLX5 VF assigned to it. This section does not explains how to create VMs either using libvirt or directly via QEMU.

Requirements

The below are the requirements for working with MLX5 VF Live Migration.

Compo nents	Description		
Adapte r Cards	 ConnectX-7 ETH BlueField-3 ETH 		
	(i) Note The same PSID must be used on both the source and the target hosts (identical cards, same CAPs and features are needed), and have the same firmware version.		
Firmwa re	 28.41.1000 32.41.1000		
Kernel	Linux v6.7 or newer		
User Space Tools	iproute2 version 6.2 or newer		
QEMU	QEMU 8.1 or newer		
Libvirt	Libvirt 8.6 or newer		

Setup

NVCONFIG

SR-IOV should be enabled and be configured to support the required number of VFs as of enabling live migration. This can be achieved by the below command:

```
mlxconfig -d *<PF_BDF>* s SRIOV_EN=1 NUM_OF_VFS=4
VF_MIGRATION_MODE=2
```

where:

SRIOV_EN	Enable Single-Root I/O Virtualization (SR-IOV)		
NUM_OF_VFS	The total number of Virtual Functions (VFs) that can be supported, for each PF.		
VF_MIGRATION_M ODE	 Defines support for VF migration. 0x0: DEVICE_DEFAULT 0x1: MIGRATION_DISABLED 0x2: MIGRATION_ENABLED 		

Kernel Configuration

Needs to be compiled with driver MLX5_VFIO_PCI enabled. (i.e. CONFIG_MLX5_VFIO_PCI).

To load the driver, run:

modprobe mlx5_vfio_pci

QEMU

Needs to be compiled with VFIO_PCI enabled (this is enabled by default).

Host Preparation

As stated earlier, creating the VMs is beyond the scope of this guide and we assume that they are already created. However, the VM configuration should be a migratable configuration, similarly to how it is done without SRIOV VFs.

(i) Note

The below steps should be done before running the VMs.

Over libvirt

1. Set the PF in the "switchdev" mode.

devlink dev eswitch set pci/<PF_BDF> mode switchdev

2. Create the VFs that will be assigned to the VMs.

echo "1" > /sys/bus/pci/devices/<PF_BDF>/sriov_numvfs

- 3. Set the VFs as migration capable.
 - 1. See the name of the VFs, run:

devlink port show

2. Unbind the VFs from mlx5_core, run:

echo '<VF_BDF>' > /sys/bus/pci/drivers/mlx5_core/unbind

3. Use devlink to set each VF as migration capable, run:

devlink port function set pci/<PF_BDF>/1 migratable
enable

- 4. Assign the VFs to the VMs.
 - 1. Edit the VMs XML file, run:

virsh edit <VM_NAME>

2. Assign the VFs to the VM by adding the following under the "devices" tag:

```
<hostdev mode='subsystem' type='pci' managed='no'>
<driver name='vfio'/>
<source>
<address domain='0x0000' bus='0x08' slot='0x00' function='0x2'/>
</source>
<address type='pci' domain='0x0000' bus='0x09' slot='0x00' function='0x0'/>
</hostdev>
```

j) Note

The domain, bus, slot and function values above are dummy values, replace them with your VFs values.

- 5. Set the destination VM in incoming mode.
 - 1. Edit the destination VM XML file, run:

```
virsh edit <VM_NAME>
```

2. Set the destination VM in migration incoming mode by adding the following under "domain" tag:

```
<domain type='kvm' xmlns:qemu='http://libvirt.org/schemas/domain/qemu/1.0'>
[...]
<qemu:commandline>
<qemu:arg value='--incoming'/>
<qemu:arg value='tcp:<DEST_IP>:<DEST_PORT>'/>
</qemu:commandline>
</domain>
```



To be able to save the file, the above "xmlns:qemu" attribute of the "domain" tag must be added as well.

6. Bind the VFs to mlx5_vfio_pci driver.

1.

1. Detach the VFs from libvirt management, run:

virsh nodedev-detach pci_<VF_BDF>

2. Unbind the VFs from vfio-pci driver (the VFs are automatically bound to it after running "virsh nodedev-detach"), run:

echo '<VF_BDF>' > /sys/bus/pci/drivers/vfio-pci/unbind

3. Set driver override, run:

```
echo 'mlx5_vfio_pci' >
/sys/bus/pci/devices/<VF_BDF>/driver_override
```

4. Bind the VFs to mlx5_vfio_pci driver, run:

echo '<VF_BDF>' > /sys/bus/pci/drivers/mlx5_vfio_pci/bind

Directly over QEMU

1. Set the PF in "switchdev" mode.

devlink dev eswitch set pci/<PF_BDF> mode switchdev

2. Create the VFs that will be assigned to the VMs.

echo "1" > /sys/bus/pci/devices/<PF_BDF>/sriov_numvfs

- 3. Set the VFs as migration capable.
 - 1. See the name of the VFs, run:

```
devlink port show
```

2. Unbind the VFs from mlx5_core, run:

```
echo '<VF_BDF>' > /sys/bus/pci/drivers/mlx5_core/unbind
```

3. Use devlink to set each VF as migration capable, run:

```
devlink port function set pci/<PF_BDF>/1 migratable
enable
```

- 4. Bind the VFs to mlx5_vfio_pci driver:
 - 1. Set driver override, run:

```
echo 'mlx5_vfio_pci' >
/sys/bus/pci/devices/<VF_BDF>/driver_override
```

2. Bind the VFs to mlx5_vfio_pci driver, run:

echo '<VF_BDF>' > /sys/bus/pci/drivers/mlx5_vfio_pci/bind

Running the Migration

Over libvirt

1. Start the VMs in source and in destination, run:

```
virsh start <VM_NAME>
```

2. Enable switchover-ack QEMU migration capability. Run the following commands both in source and destination:

```
virsh qemu-monitor-command <VM_NAME> --hmp "migrate_set_capability
return-path on"
```

virsh qemu-monitor-command <VM_NAME> --hmp "migrate_set_capability
switchover-ack on"

3. **[Optional]** Configure the migration bandwidth and downtime limit in source side:

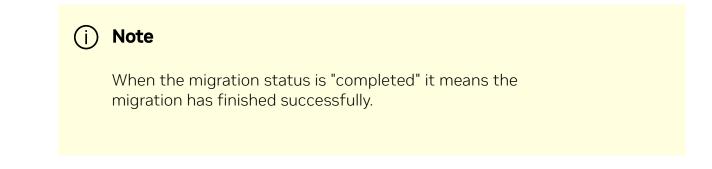
```
virsh qemu-monitor-command <VM_NAME> --hmp "migrate_set_parameter max-
bandwidth <VALUE>"
virsh qemu-monitor-command <VM_NAME> --hmp "migrate_set_parameter
```

- downtime-limit <VALUE>"
- 4. Start migration by running the migration command in source side:

```
virsh qemu-monitor-command <VM_NAME> --hmp "migrate -d tcp:<DEST_IP>:
<DEST_PORT>"
```

5. Check the migration status by running the info command in source side:

```
virsh qemu-monitor-command <VM_NAME> --hmp "info migrate"
```



Directly over QEMU

1. Start the VM in source with the VF assigned to it:

```
qemu-system-x86_64 [...] -device vfio-pci,host=
<VF_BDF>,id=mlx5_1
```

2. Start the VM in destination with the VF assigned to it and with the "incoming" parameter:

```
qemu-system-x86_64 [...] -device vfio-pci,host=
<VF_BDF>,id=mlx5_1 -incoming tcp:<DEST_IP>:<DEST_PORT>
```

3. Enable switchover-ack QEMU migration capability. Run the following commands in QEMU monitor, both in source and destination:

```
migrate_set_capability return-path on
```

```
migrate_set_capability switchover-ack on
```

4. **[Optional]** Configure the migration bandwidth and downtime limit in source side:

```
migrate_set_parameter max-bandwidth <VALUE>
migrate_set_parameter downtime-limit <VALUE>
```

5. Start migration by running the migration command in QEMU monitor in source side:

```
migrate -d tcp:<DEST_IP>:<DEST_PORT>
```

6. Check the migration status by running the info command in QEMU monitor in source side:

info migrate



When the migration status is "completed" it means the migration has finished successfully.

Migration with MultiPort vHCA

Enables the usage of a dual port Virtual HCA (vHCA) to share RDMA resources (e.g., MR, CQ, SRQ, PDs) across the two Ethernet (RoCE) NIC network ports and display the NIC as a dual port device.

MultiPort vHCA (MPV) VF is made of 2 "regular" VFs, one VF of each port. Creating a migratable MPV VF requires the same steps as regular VF (see steps in section <u>Over</u> <u>libvirt</u>). The steps should be performed on each of the NIC ports. MPV VFs traffic cannot be configured with OVS. TC rules must be defined to configure the MPV VFs traffic.

Notes

(i) Note

In ConnectX-7 adapter cards, migration cannot run in parallel on more than 4 VFs. It is the administrator's responsibility to control that.

i Note

Live migration requires same firmware version on both the source and the target hosts.

Enabling Paravirtualization

To enable Paravirtualization:



The example below works on RHEL7.* without a Network Manager.

1. Create a bridge.

```
vim /etc/sysconfig/network-scripts/ifcfg-bridge0
DEVICE=bridge0
TYPE=Bridge
IPADDR=12.195.15.1
NETMASK=255.255.0.0
B00TPROT0=static
ONB00T=yes
```

```
NM_CONTROLLED=no
DELAY=0
```

2. Change the related interface (in the example below bridge0 is created over eth5).

DEVICE=eth5 BOOTPROTO=none STARTMODE=on HWADDR=00:02:c9:2e:66:52 TYPE=Ethernet NM_CONTROLLED=no ONBOOT=yes BRIDGE=bridge0

- 3. Restart the service network.
- 4. Attach a bridge to VM.

```
ifconfig -a
...
eth6 Link encap:Ethernet HWaddr 52:54:00:E7:77:99
    inet addr:13.195.15.5 Bcast:13.195.255.255 Mask:255.255.00
    inet6 addr: fe80::5054:ff:fee7:7799/64 Scope:Link
    UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
    RX packets:481 errors:0 dropped:0 overruns:0 frame:0
    TX packets:450 errors:0 dropped:0 overruns:0
carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:22440 (21.9 KiB) TX bytes:19232 (18.7 KiB)
    Interrupt:10 Base address:0xa000
...
```

VXLAN Hardware Stateless Offloads

VXLAN technology provides scalability and security challenges solutions. It requires extension of the traditional stateless offloads to avoid performance drop. ConnectX family cards offer the following stateless offloads for a VXLAN packet, similar to the ones offered to non-encapsulated packets. VXLAN protocol encapsulates its packets using outer UDP header.

Available hardware stateless offloads:

- Checksum generation (Inner IP and Inner TCP/UDP)
- Checksum validation (Inner IP and Inner TCP/UDP)
- TSO support for inner TCP packets
- RSS distribution according to inner packets attributes
- Receive queue selection inner frames may be steered to specific QPs

Enabling VXLAN Hardware Stateless Offloads

VXLAN offload is enabled by default for ConnectX-4 family devices running the minimum required firmware version and a kernel version that includes VXLAN support.

To confirm if the current setup supports VXLAN, run:

```
ethtool -k $DEV | grep udp_tnl
```

Example:

```
ethtool -k ens1f0 | grep udp_tnl
tx-udp_tnl-segmentation: on
```

ConnectX-4 family devices support configuring multiple UDP ports for VXLAN offload. Ports can be added to the device by configuring a VXLAN device from the OS command line using the "ip" command.

Note: If you configure multiple UDP ports for offload and exceed the total number of ports supported by hardware, then those additional ports will still function properly, but will not benefit from any of the stateless offloads.

Example:

```
ip link add vxlan0 type vxlan id 10 group 239.0.0.10 ttl 10 dev
ens1f0 dstport 4789
ip addr add 192.168.4.7/24 dev vxlan0
ip link set up vxlan0
```

Note: dstport' parameters are not supported in Ubuntu 14.4.

The VXLAN ports can be removed by deleting the VXLAN interfaces.

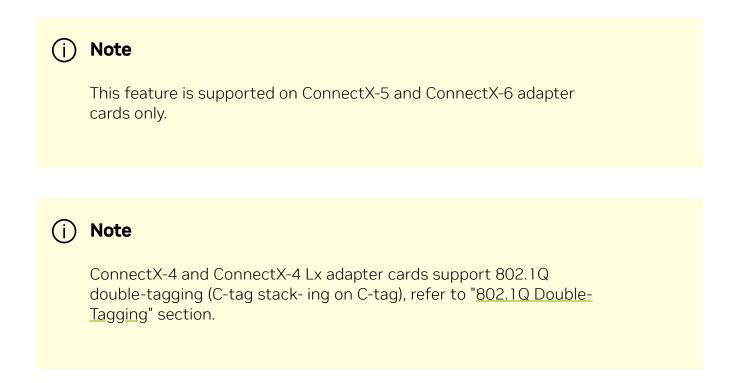
Example:

ip link delete vxlan0

Important Note

VXLAN tunneling adds 50 bytes (14-eth + 20-ip + 8-udp + 8-vxlan) to the VM Ethernet frame. Please verify that either the MTU of the NIC who sends the packets, e.g. the VM virtio-net NIC or the host side veth device or the uplink takes into account the tunneling overhead. Meaning, the MTU of the sending NIC has to be decremented by 50 bytes (e.g 1450 instead of 1500), or the uplink NIC MTU has to be incremented by 50 bytes (e.g 1550 instead of 1500)

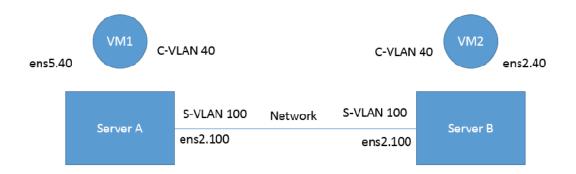
Q-in-Q Encapsulation per VF in Linux (VST)



This section describes the configuration of IEEE 802.1ad QinQ VLAN tag (S-VLAN) to the hypervisor per Virtual Function (VF). The Virtual Machine (VM) attached to the VF (via SR-IOV) can send traffic with or without C-VLAN. Once a VF is configured to VST QinQ encapsulation (VST QinQ), the adapter's hardware will insert S-VLAN to any packet from the VF to the physical port. On the receive side, the adapter hardware will strip the S-VLAN from any packet coming from the wire to that VF.

Setup

The setup assumes there are two servers equipped with ConnectX-5/ConnectX-6 adapter cards.



Prerequisites

- Kernel must be of v3.10 or higher, or custom/inbox kernel must support vlan-stag
- Firmware version 16/20.21.0458 or higher must be installed for ConnectX-5/ConnectX-6 HCAs
- The server should be enabled in SR-IOV and the VF should be attached to a VM on the hypervisor.
 - In order to configure SR-IOV in Ethernet mode for ConnectX-5/ConnectX-6 adapter cards, please refer to "<u>Configuring SR-IOV for ConnectX-4/ConnectX-5</u> <u>(Ethernet)</u>" section. In the following configuration example, the VM is attached to VF0.
- Network Considerations the network switches may require increasing the MTU (to support 1522 MTU size) on the relevant switch ports.

Configuring Q-in-Q Encapsulation per Virtual Function for ConnectX-5/ConnectX-6

- 1. Add the required S-VLAN (QinQ) tag (on the hypervisor) per port per VF. There are two ways to add the S-VLAN:
 - 1. By using sysfs:

echo '100:0:802.1ad' > /sys/class/net/ens1f0/device/sriov/0/vlan

2. By using the ip link command (available only when using the latest Kernel version):

ip link set dev ens1f0 vf 0 vlan 100 proto 802.1ad

Check the configuration using the ip link show command:

```
# ip link show ens1f0
ens1f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc
mq state UP mode DEFAULT qlen 1000
```

```
link/ether ec:0d:9a:44:37:84 brd ff:ff:ff:ff:ff:ff
vf 0 MAC 00:00:00:00:00, vlan 100, vlan protocol
802.1ad, spoof checking off, link-state auto, trust off
vf 1 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 2 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 3 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 4 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
```

2. **Optional:** Add S-VLAN priority. Use the qos parameter in the ip link command (or sysfs):

ip link set dev ens1f0 vf 0 vlan 100 qos 3 proto 802.1ad

Check the configuration using the ip link show command:

```
# ip link show ens1f0
ens1f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq
state UP mode DEFAULT qlen 1000
link/ether ec:0d:9a:44:37:84 brd ff:ff:ff:ff:ff
vf 0 MAC 00:00:00:00:00, vlan 100, qos 3, vlan protocol
802.1ad, spoof checking off, link-state auto, trust off
vf 1 MAC 00:00:00:00:00, spoof checking off, link-state
auto, trust off
vf 2 MAC 00:00:00:00:00, spoof checking off, link-state
auto, trust off
vf 3 MAC 00:00:00:00:00, spoof checking off, link-state
auto, trust off
```

```
vf 4 MAC 00:00:00:00:00, spoof checking off, link-state
auto, trust off
```

3. Create a VLAN interface on the VM and add an IP address.

```
ip link add link ens5 ens5.40 type vlan protocol 802.1q id 40
ip addr add 42.134.135.7/16 brd 42.134.255.255 dev ens5.40
ip link set dev ens5.40 up
```

4. To verify the setup, run ping between the two VMs and open Wireshark or tcpdump to capture the packet.

802.1Q Double-Tagging

This section describes the configuration of 802.1Q double-tagging support to the hypervisor per Virtual Function (VF). The Virtual Machine (VM) attached to the VF (via SR-IOV) can send traffic with or without C-VLAN. Once a VF is configured to VST encapsulation, the adapter's hardware will insert C-VLAN to any packet from the VF to the physical port. On the receive side, the adapter hardware will strip the C-VLAN from any packet coming from the wire to that VF.

Configuring 802.1Q Double-Tagging per Virtual Function

- 1. Add the required C-VLAN tag (on the hypervisor) per port per VF. There are two ways to add the C-VLAN:
 - 1. By using sysfs:

echo '100:0:802.1q' > /sys/class/net/ens1f0/device/sriov/0/vlan

2. By using the ip link command (available only when using the latest Kernel version):

ip link set dev ens1f0 vf 0 vlan 100

Check the configuration using the ip link show command:

```
# ip link show ens1f0
ens1f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc
mq state UP mode DEFAULT qlen 1000
link/ether ec:0d:9a:44:37:84 brd ff:ff:ff:ff:ff
vf 0 MAC 00:00:00:00:00, vlan 100, spoof checking
off, link-state auto, trust off
vf 1 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 2 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 3 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
vf 4 MAC 00:00:00:00:00, spoof checking off, link-
state auto, trust off
```

2. Create a VLAN interface on the VM and add an IP address.

```
# ip link add link ens5 ens5.40 type vlan protocol 802.1q id
40
# ip addr add 42.134.135.7/16 brd 42.134.255.255 dev ens5.40
# ip link set dev ens5.40 up
```

3. To verify the setup, run ping between the two VMs and open Wireshark or tcpdump to capture the packet.

Scalable Functions

Scalable function is a lightweight function that has a parent PCI function on which it is deployed. Scalable functions are useful for containers where netdevice and RDMA devices of a scalable function can be assigned to a container. This way, the container can get complete offload capabilities of an eswitch, isolation and dedicated accelerated network device. For Step-by-Step Configuration instructions, follow the User Guide <u>here</u>.

Resiliency

The chapter contains the following sections:

• <u>Reset Flow</u>

Reset Flow

Reset Flow is activated by default. Once a "fatal device" error is recognized, both the HCA and the software are reset, the ULPs and user application are notified about it, and a recovery process is performed once the event is raised.

Currently, a reset flow can be triggered by a firmware assert with Recover Flow Request (RFR) only. Firmware RFR support should be enabled explicitly using mlxconfig commands.

To query the current value, run:

mlxconfig -d /dev/mst/mt4115_pciconf0 query | grep SW_RECOVERY_ON_ERRORS

To enable RFR bit support, run:

mlxconfig -d /dev/mst/mt4115_pciconf0 set SW_RECOVERY_ON_ERRORS=true

Kernel ULPs

Once a "fatal device" error is recognized, an IB_EVENT_DEVICE_FATAL event is created, ULPs are notified about the incident, and outstanding WQEs are simulated to be returned with "flush in error" message to enable each ULP to close its resources and not get stuck via calling its "remove_one" callback as part of "Reset Flow".

Once the unload part is terminated, each ULP is called with its "add_one" callback, its resources are re-initialized and it is re-activated.

User Space Applications (IB/RoCE)

Once a "fatal device" error is recognized an IB_EVENT_DEVICE_FATAL event is created, applications are notified about the incident and relevant recovery actions are taken.

Applications that ignore this event enter a zombie state, where each command sent to the kernel is returned with an error, and no completion on outstanding WQEs is expected.

The expected behavior from the applications is to register to receive such events and recover once the above event is raised. Same behavior is expected in case the NIC is unbounded from the PCI and later is rebounded. Applications running over RDMA CM should behave in the same manner once the RDMA_CM_EVENT_DEVICE_REMOVAL event is raised.

The below is an example of using the unbind/bind for NIC defined by "0000:04:00.0"

echo 0000:04:00.0 > /sys/bus/pci/drivers/mlx5_core/unbind echo 0000:04:00.0 > /sys/bus/pci/drivers/mlx5_core/bind

SR-IOV

If the Physical Function recognizes the error, it notifies all the VFs about it by marking their communication channel with that information, consequently, all the VFs and the PF are reset.

If the VF encounters an error, only that VF is reset, whereas the PF and other VFs continue to work unaffected.

Forcing the VF to Reset

If an outside "reset" is forced by using the PCI sysfs entry for a VF, a reset is executed on that VF once it runs any command over its communication channel.

For example, the below command can be used on a hypervisor to reset a VF defined by 0000:04:00.1:

```
echo 1 >/sys/bus/pci/devices/0000:04:00.1/reset
```

Extended Error Handling (EEH)

Extended Error Handling (EEH) is a PowerPC mechanism that encapsulates AER, thus exposing AER events to the operating system as EEH events.

The behavior of ULPs and user space applications is identical to the behavior of AER.

CRDUMP

CRDUMP feature allows for taking an automatic snapshot of the device CR-Space in case the device's FW/HW fails to function properly.

Snapshots Triggers:

The snapshot is triggered after firmware detects a critical issue, requiring a recovery flow.

This snapshot can later be investigated and analyzed to track the root cause of the failure.

Currently, only the first snapshot is stored, and is exposed using a temporary virtual file. The virtual file is cleared upon driver reset.

When a critical event is detected, a message indicating CRDUMP collection will be printed to the Linux log. User should then back up the file pointed to in the printed message. The file location format is: /proc/driver/mlx5_core/crdump/<pci address>

Snapshot should be copied by Linux standard tool for future investigation.

Firmware Tracer

This mechanism allows for the device's FW/HW to log important events into the event tracing system (/sys/kernel/debug/tracing) without requiring any NVIDIA tool.



To be able to use this feature, trace points must be enabled in the kernel.

This feature is enabled by default, and can be controlled using sysfs commands.

To disable the feature:

echo 0 > /sys/kernel/debug/tracing/events/mlx5/fw_tracer/enable

To enable the feature:

echo 1 > /sys/kernel/debug/tracing/events/mlx5/fw_tracer/enable

To view FW traces using vim text editor:

vim /sys/kernel/debug/tracing/trace

Docker Containers

On Linux, Docker uses resource isolation of the Linux kernel, to allow independent "containers" to run within a single Linux kernel instance.

Docker containers are supported on MLNX_OFEDusing Docker runtime. Virtual RoCE and InfiniBand devices are supported using SR-IOV mode.

Currently, RDMA/RoCE devices are supported in the modes listed in the following table.

Linux Containers Networking Modes

Orchestration and Clustering Tool	Version	Networking Mode	Link Layer	Virtualizat ion Mode
Docker	Docker Engine 17.03 or higher	SR-IOV using sriov-plugin along with docker run wrapper tool	InfiniBand and Ethernet	SR-IOV
Kubernetes	Kubernet es 1.10.3 or higher	SR-IOV using device plugin, and using SR- IOV CNI plugin	InfiniBand and Ethernet	SR-IOV
		VXLAN using IPoIB bridge	InfiniBand	Shared HCA

Docker Using SR-IOV

In this mode, Docker engine is used to run containers along with SR-IOV networking plugin. To isolate the virtual devices, docker_rdma_sriov tool should be used. This mode is applicable to both InfiniBand and Ethernet link layers.

To obtain the plugin, visit: <u>hub.docker.com/r/rdma/sriov-plugin</u>

To install the docker_rdma_sriov tool, use the container tools installer available via <u>hub.docker.com/r/rdma/container_tools_installer</u>

For instructions on how to use Docker with SR-IOV, refer to <u>Docker RDMA SRIOV</u> <u>Networking with ConnectX4/ConnectX5/ConnectX6</u> Community post.

Kubernetes Using SR-IOV

In order to use RDMA in Kubernetes environment with SR-IOV networking mode, two main components are required:

- 1. RDMA device plugin this plugin allows for exposing RDMA devices in a Pod
- 2. SR-IOV CNI plugin this plugin provisions VF net device in a Pod

When used in SR-IOV mode, this plugin enables SR-IOV and performs necessary configuration including setting GUID, MAC, privilege mode, and Trust mode.

The plugin also allocates the VF devices when Pods are scheduled and requested by Kubernetes framework.

Kubernetes with Shared HCA

One RDMA device (HCA) can be shared among multiple Pods running in a Kubernetes worker nodes. User defined networks are created using VXLAN or VETH networking devices. RDMA device (HCA) can be shared among multiple Pods running in a Kubernetes worker nodes.

HPC-X

For information on HPC-X®, please refer to HPC-X User Manual at <u>developer.nvidia.com/networking/hpc-x</u>.

Fast Driver Unload

This feature enables optimizing mlx5 driver teardown time in shutdown and kexec flows.

The fast driver unload is disabled by default. To enable it, the prof_sel module parameter of mlx5_core module should be set to 3.

OVS Offload Using ASAP² Direct

Overview



Supported on ConnectX-5 and above adapter cards.

Open vSwitch (OVS) allows Virtual Machines (VMs) to communicate with each other and with the outside world. OVS traditionally resides in the hypervisor and switching is based

on twelve tuple matching on flows. The OVS software based solution is CPU intensive, affecting system performance and preventing full utilization of the available bandwidth.

NVIDIA Accelerated Switching And Packet Processing (ASAP²) technology allows OVS offloading by handling OVS data-plane in ConnectX-5 onwards NIC hardware (Embedded Switch or eSwitch) while maintaining OVS control-plane unmodified. As a result, we observe significantly higher OVS performance without the associated CPU load.

As of v5.0, OVS-DPDK became part of MLNX_OFED package. OVS-DPDK supports $ASAP^2$ just as the OVS-Kernel (Traffic Control (TC) kernel-based solution) does, yet with a different set of features.

The traditional ASAP² hardware data plane is built over SR-IOV virtual functions (VFs), so that the VF is passed through directly to the VM, with the NVIDIA driver running within the VM. An alternate approach that is also supported is vDPA (vhost Data Path Acceleration). vDPA allows the connection to the VM to be established using VirtIO, so that the data-plane is built between the SR-IOV VF and the standard VirtIO driver within the VM, while the control-plane is managed on the host by the vDPA application. Two flavors of vDPA are supported, Software vDPA; and Hardware vDPA. Software vDPA management functionality is embedded into OVS-DPDK, while Hardware vDPA uses a standalone application for management, and can be run with both OVS-Kernel and OVS-DPDK. For further information, please see sections <u>VirtIO Acceleration through VF Relay</u> (Software vDPA) and <u>VirtIO Acceleration through Hardware vDPA</u>.

Installing OVS-Kernel ASAP² Packages

Install the required packages. For the complete solution, you need to install supporting MLNX_OFED(v4.4 and above), iproute2, and openvswitch packages.

Installing OVS-DPDK ASAP² Packages

Run:

./mlnxofedinstall --ovs-dpdk -upstream-libs

Setting Up SR-IOV

(i) Note

Note that this section applies to both OVS-DPDK and OVS-Kernel similarly.

To set up SR-IOV:

1. Choose the desired card.

The example below shows a dual-ported ConnectX-5 card (device ID 0x1017) and **a** single SR-IOV VF (Virtual Function, device ID 0x1018).

In SR-IOV terms, the card itself is referred to as the PF (Physical Function).

```
# lspci -nn | grep Mellanox
Oa:00.0 Ethernet controller [0200]: Mellanox Technologies
MT27800 Family [ConnectX-5] [15b3:1017]
Oa:00.1 Ethernet controller [0200]: Mellanox Technologies
MT27800 Family [ConnectX-5] [15b3:1017]
Oa:00.2 Ethernet controller [0200]: Mellanox Technologies
MT27800 Family [ConnectX-5 Virtual Function] [15b3:1018]
```



Enabling SR-IOV and creating VFs is done by the firmware upon admin directive as explained in Step 5 below.

2. Identify the NVIDIA NICs and locate net-devices which are on the NIC PCI BDF.

ls -l /sys/class/net/ | grep 04:00

lrwxrwxrwx 1 root root 0 Mar 27 16:58 enp4s0f0 ->
../../devices/pci0000:00/0000:00:03.0/0000:04:00.0/net/enp4s0f0
lrwxrwxrwx 1 root root 0 Mar 27 16:58 enp4s0f1 ->
../../devices/pci0000:00/0000:00:03.0/0000:04:00.1/net/enp4s0f1
lrwxrwxrwx 1 root root 0 Mar 27 16:58 eth0 ->
../../devices/pci0000:00/0000:00:03.0/0000:04:00.2/net/eth0
lrwxrwxrwx 1 root root 0 Mar 27 16:58 eth1 ->
../../devices/pci0000:00/0000:00:03.0/0000:04:00.3/net/eth1

The PF NIC for port #1 is enp4s0f0, and the rest of the commands will be issued on it.

3. Check the firmware version.

Make sure the firmware versions installed are as state in the Release Notes document.

```
# ethtool -i enp4s0f0 | head -5
driver: mlx5_core
version: 5.0-5
firmware-version: 16.21.0338
expansion-rom-version:
bus-info: 0000:04:00.0
```

4. Make sure SR-IOV is enabled on the system (server, card).

Make sure SR-IOV is enabled by the server BIOS, and by the firmware with up to N VFs, where N is the number of VFs required for your environment. Refer to "<u>NVIDIA</u> <u>Firmware Tools</u>" below for more details.

```
# cat /sys/class/net/enp4s0f0/device/sriov_totalvfs
4
```

5. Turn ON SR-IOV on the PF device.

echo 2 > /sys/class/net/enp4s0f0/device/sriov_numvfs

6. Provision the VF MAC addresses using the IP tool.

ip link set enp4s0f0 vf 0 mac e4:11:22:33:44:50
ip link set enp4s0f0 vf 1 mac e4:11:22:33:44:51

7. Verify the VF MAC addresses were provisioned correctly and SR-IOV was turned ON.

In the example above, the maximum number of possible VFs supported by the firmware is 4 and only 2 are enabled.

8. Provision the PCI VF devices to VMs using PCI Pass-Through or any other preferred virt tool of choice, e.g virt-manager.

For further information on SR-IOV, refer to <u>HowTo Configure SR-IOV for ConnectX-4/ConnectX-5/ConnectX-6 with KVM (Ethernet)</u>.

OVS Hardware Offloads Configuration

OVS-Kernel Hardware Offloads

SwitchDev Configuration

1. Unbind the VFs.

```
echo 0000:04:00.2 > /sys/bus/pci/drivers/mlx5_core/unbind
echo 0000:04:00.3 > /sys/bus/pci/drivers/mlx5_core/unbind
```



VMs with attached VFs must be powered off to be able to unbind the VFs.

2. Change the eSwitch mode from Legacy to SwitchDev on the PF device.

This will also create the VF representor netdevices in the host OS.

devlink dev eswitch set pci/0000:3b:00.0 mode switchdev



Before changing the mode, make sure that all VFs are unbound.



To go back to SR-IOV legacy mode, run:

devlink dev eswitch set pci/0000:3b:00.0 mode legacy

This will also remove the VF representor netdevices.

On old OSs or kernels that do not support Devlink, moving to SwitchDev mode can be done using sysfs.

echo switchdev > /sys/class/net/enp4s0f0/compat/devlink/mode

3. At this stage, VF representors have been created. To map representor to its VF, make sure to obtain the representor's switchid and portname from:

switchid - used to map representor to device, both device PFs have the same switchid.

portname - used to map representor to PF and VF, value returned is pfXvfY, where X is the PF number and Y is the number of VF.

On old kernels, switchid and portname can be acquired through sysfs:

4. Bind the VFs.

echo 0000:04:00.2 > /sys/bus/pci/drivers/mlx5_core/bind echo 0000:04:00.3 > /sys/bus/pci/drivers/mlx5_core/bind

SwitchDev Performance Tuning

SwitchDev performance can be further improved by tuning it.

Steering Mode

OVS-kernel supports two steering modes for rules insertion into hardware.

1. **SMFS** – Software Managed Flow Steering (as of MLNX_OFED v5.1, this is the default mode)

Rules are inserted directly to the hardware by the software (driver). This mode is optimized for rules insertion.

2. DMFS - Device Managed Flow Steering

Rules insertion is done using firmware commands. This mode is optimized for throughput with a small amount of rules in the system.

The mode can be controlled via sysfs or devlink API in kernels that support it:

```
Sysfs:
# echo smfs > /sys/class/net/<PF
netdev>/compat/devlink/steering_mode
Devlink:
# devlink dev param set pci/0000:00:08.0 name
flow_steering_mode value "smfs" cmode runtime
Replace smfs param with dmfs for device managed flow steering
```

Notes:

- The mode should be set prior to moving to SwitchDev, by echoing to the sysfs or invoking the devlink command.
- Only when moving to SwitchDev will the driver use the mode set by the previous step.

- Mode cannot be changed after moving to SwitchDev.
- The steering mode is applicable for SwitchDev mode only, meaning it does not affect legacy SR-IOV or other configurations.

Troubleshooting SMFS

mlx5 debugfs was extended to support presenting Software Steering resources:

dr_domain including it's tables, matchers and rules.

The interface is read-only.

While dump is being created, new steering rules cannot be inserted/deleted.

The steering information is dumped in the CSV form with the following format: <object_type>,<object_ID>, <object_info>,...,<object_info>

This data can be read at the following path: /sys/kernel/debug/mlx5/<BDF>/steering/fdb/<domain_handle>

Example:

cat
/sys/kernel/debug/mlx5/0000:82:00.0/steering/fdb/dmn_000018644
3100,0x55caa4621c50,0xee802,4,65533
3101,0x55caa4621c50,0xe0100008

You can then use the steering dump parser to make the output more human readable.

The parser can be found in the following public GitHub repository: <u>https://github.com/Mellanox/mlx_steering_dump</u>

vPort Match Mode

OVS-kernel support two modes that define how the rules on match on vport.

1. Metadata - rules match on metadata instead of vport number (default mode).

This mode is needed in order to support SR-IOV Live migration and Dual port RoCE features.

Matching on Metadata can have a performance impact.

2. Legacy – rules match on vport number.

In this mode, performance can be higher in comparison to Metadata. It can still be used only if none of the above features (SR-IOV Live migration and Dual port RoCE) is enabled/used.

The mode can be controlled via sysfs:

```
Set Legacy:
# echo legacy > /sys/class/net/<PF
netdev>/compat/devlink/vport_match_mode
Set metadata:
Devlink:
# echo metadata > /sys/class/net/<PF
netdev>/compat/devlink/vport_match_mode
```

Note: This mode should be set prior to moving to SwitchDev, by echoing to the sysfs.

Flow Table Large Group Number

Offloaded flows, including Connection Tracking, are added to Virtual Switch Forwarding Data Base (FDB) flow tables. FDB tables have a set of flow groups, where each flow group saves the same traffic pattern flows. E.g, for connection tracking offloaded flow, TCP and UDP are different traffic patterns which will end up in two different flow groups.

A flow group has a limited size to save flow entries. As default, the driver has 15 big FDB flow groups. Each of these big flow groups can save 4M / (15 + 1) = 256k different 5-tuple flow entries at most. For scenarios with more than 15 traffic patterns, the driver provides a module parameter (num_of_groups) to allow customization and performance tuning.

The mode can be controlled via module param or devlink API for kernels that support it:

Notes:

- In MLNX_OFED v5.1, the default value was changed from 4 to 15.
- The change takes effect immediately if there is no flow inside the FDB table (no traffic running and all offloaded flows are aged out). And it can be dynamically changed without reloading the driver.

If there are still offloaded flows residual when changing this parameter, it will only take effect after all flows have aged out.

Open vSwitch Configuration

Open vSwitch configuration is a simple OVS bridge configuration with SwitchDev.

1. Run the openvswitch service.

systemctl start openvswitch

2. Create an OVS bridge (here it's named ovs-sriov).

```
# ovs-vsctl add-br ovs-sriov
```

3. Enable hardware offload (disabled by default).

```
# ovs-vsctl set Open_vSwitch . other_config:hw-offload=true
```

4. Restart the openvswitch service. This step is required for HW offload changes to take effect.

systemctl restart openvswitch



HW offload policy can also be changed by setting the tc-policy using one on the following values:

* none - adds a TC rule to both the software and the hardware (default)

* skip_sw - adds a TC rule only to the hardware

* skip_hw - adds a TC rule only to the software

The above change is used for debug purposes.

5. Add the PF and the VF representor netdevices as OVS ports.

ovs-vsctl add-port ovs-sriov enp4s0f0

```
# ovs-vsctl add-port ovs-sriov enp4s0f0_0
```

ovs-vsctl add-port ovs-sriov enp4s0f0_1

Make sure to bring up the PF and representor netdevices.

ip link set dev enp4s0f0 up
ip link set dev enp4s0f0_0 up
ip link set dev enp4s0f0_1 up

The PF represents the uplink (wire).

```
# ovs-dpctl show
system@ovs-system:
    lookups: hit:0 missed:192 lost:1
    flows: 2
    masks: hit:384 total:2 hit/pkt:2.00
    port 0: ovs-system (internal)
    port 1: ovs-sriov (internal)
    port 2: enp4s0f0
    port 3: enp4s0f0_0
    port 4: enp4s0f0_1
```

6. Run traffic from the VFs and observe the rules added to the OVS data-path.

```
# ovs-dpctl dump-flows
recirc_id(0),in_port(3),eth(src=e4:11:22:33:44:50,dst=e4:1d:2d:a5
eth_type(0x0800),ipv4(frag=no), packets:33, bytes:3234,
used:1.196s, actions:2
recirc_id(0),in_port(2),eth(src=e4:1d:2d:a5:f3:9d,dst=e4:11:22:3
```

```
eth_type(0x0800),ipv4(frag=no), packets:34, bytes:3332,
used:1.196s, actions:3
```

In the example above, the ping was initiated from VFO (OVS port 3) to the outer node (OVS port 2), where the VF MAC is e4:11:22:33:44:50 and the outer node MAC is e4:11:22:a5:f3:9d

As shown above, two OVS rules were added, one in each direction.

Note that you can also verify offloaded packets by adding type=offloaded to the command. For example:

ovs-appctl dpctl/dump-flows type=offloaded

Open vSwitch Performance Tuning

Flow Aging

The aging timeout of OVS is given in ms and can be controlled using the following command.

ovs-vsctl set Open_vSwitch . other_config:max-idle=30000

TC Policy

Specifies the policy used with HW offloading.

- none adds a TC rule to both the software and the hardware (default)
- skip_sw adds a TC rule only to the hardware
- **skip_hw** adds a TC rule only to the software

Example:

```
# ovs-vsctl set Open_vSwitch . other_config:tc-policy=skip_sw
```

Note: TC policy should only be used for debugging purposes.

Max-Revalidator

Specifies the maximum time (in ms) that revalidator threads will wait for kernel statistics before executing flow revalidation.

ovs-vsctl set Open_vSwitch . other_config:max-revalidator=10000

n-handler-threads

Specifies the number of threads for software datapaths to use for handling new flows.

The default value is the number of online CPU cores minus the number of revalidators.

```
# ovs-vsctl set Open_vSwitch . other_config:n-handler-threads=4
```

n-revalidator-threads

Specifies the number of threads for software datapaths to use for revalidating flows in the datapath.

```
# ovs-vsctl set Open_vSwitch . other_config:n-revalidator-
threads=4
```

vlan-limit

Limits the number of VLAN headers that can be matched to the specified number.

```
# ovs-vsctl set Open_vSwitch . other_config:vlan-limit=2
```

Basic TC Rules Configuration

Offloading rules can also be added directly, and not only through OVS, using the tc utility.

To create an offloading rule using TC:

1. Create an ingress qdisc (queueing discipline) for each interface that you wish to add rules into.

tc qdisc add dev enp4s0f0 ingress
tc qdisc add dev enp4s0f0_0 ingress
tc qdisc add dev enp4s0f0_1 ingress

2. Add TC rules using flower classifier in the following format.

```
# tc filter add dev NETDEVICE ingress protocol PROTOCOL prio
PRIORITY \
[chain CHAIN] flower [ MATCH_LIST ] [ action ACTION_SPEC ]
```

Note: List of supported matches (specifications) and actions can be found in <u>Classification Fields (Matches)</u> section.

3. Dump the existing tc rules using flower classifier in the following format.

tc [-s] filter show dev NETDEVICE ingress

SR-IOV VF LAG

SR-IOV VF LAG allows the NIC's physical functions (PFs) to get the rules that the OVS will try to offload to the bond net-device, and to offload them to the hardware e-switch. Bond modes supported are:

- Active-Backup
- XOR
- LACP

SR-IOV VF LAG enables complete offload of the LAG functionality to the hardware. The bonding creates a single bonded PF port. Packets from up-link can arrive from any of the physical ports, and will be forwarded to the bond device.

When hardware offload is used, packets from both ports can be forwarded to any of the VFs. Traffic from the VF can be forwarded to both ports according to the bonding state. Meaning, when in active-backup mode, only one PF is up, and traffic from any VF will go through this PF. When in XOR or LACP mode, if both PFs are up, traffic from any VF will split between these two PFs.

SR-IOV VF LAG Configuration on ASAP2

To enable SR-IOV VF LAG, both physical functions of the NIC should first be configured to SR-IOV SwitchDev mode, and only afterwards bond the up-link representors.

The example below shows the creation of bond interface on two PFs:

1. Load bonding device and enslave the up-link representor (currently PF) net-device devices.

modprobe bonding mode=802.3ad
Ifup bond0 (make sure ifcfg file is present with desired bond
configuration)

```
ip link set enp4s0f0 master bond0
ip link set enp4s0f1 master bond0
```

2. Add the VF representor net-devices as OVS ports. If tunneling is not used, add the bond device as well.

```
ovs-vsctl add-port ovs-sriov bond0
ovs-vsctl add-port ovs-sriov enp4s0f0_0
ovs-vsctl add-port ovs-sriov enp4s0f1_0
```

3. Make sure to bring up the PF and the representor netdevices.

ip link set dev bond0 up
ip link set dev enp4s0f0_0 up
ip link set dev enp4s0f1_0 up

(i) Note

Once SR-IOV VF LAG is configured, all VFs of the two PFs will become part of the bond, and will behave as described above.

Limitations

- 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.
- VF LAG configuration is not supported when the NUM_OF_VFS configured in mlxconfig is higher than 64.

Using TC with VF LAG

Both rules can be added using either of the following.

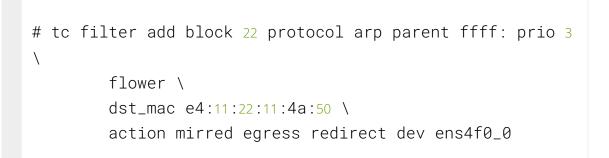
1. Shared block (supported from kernel 4.16 and RHEL/CentOS 7.7 and above).

tc qdisc add dev bond0 ingress_block 22 ingress # tc qdisc add dev ens4p0 ingress_block 22 ingress # tc qdisc add dev ens4p1 ingress_block 22 ingress

1. Add drop rule.

```
# tc filter add block 22 protocol arp parent ffff: prio 3
\
    flower \
    dst_mac e4:11:22:11:4a:51 \
    action drop
```

2. Add redirect rule from bond to representor.



3. Add redirect rule from representor to bond.

```
# tc filter add dev ens4f0_0 protocol arp parent ffff:
prio 3 \
    flower \
        dst_mac ec:0d:9a:8a:28:42 \
        action mirred egress redirect dev bond0
```

- 2. Without shared block (supported from kernel 4.15 and **below**).
 - 1. Add redirect rule from bond to representor.

```
# tc filter add dev bond0 protocol arp parent ffff: prio
1 \
        flower \
        dst_mac e4:11:22:11:4a:50 \
        action mirred egress redirect dev ens4f0_0
```

2. Add redirect rule from representor to bond.

```
# tc filter add dev ens4f0_0 protocol arp parent ffff:
prio 3 \
        flower \
        dst_mac ec:0d:9a:8a:28:42 \
        action mirred egress redirect dev bond0
```

Classification Fields (Matches)

OVS-Kernel supports multiple classification fields which packets can fully or partially match.

Ethernet Layer 2

- Destination MAC
- Source MAC
- Ethertype

Supported on all kernels.

In OVS dump flows:

```
skb_priority(0/0), skb_mark(0/0), in_port(eth6), eth(src=00:02:10:40:10:{
,dst=68:54:ed:00:af:de), eth_type(0x8100), packets:1981, bytes:206024,
used:0.440s, dp:tc, actions:eth7
```

Using TC rules:

```
tc filter add dev $rep parent ffff: protocol arp pref 1 \
flower \
dst_mac e4:1d:2d:5d:25:35 \
src_mac e4:1d:2d:5d:25:34 \
action mirred egress redirect dev $NIC
```

IPv4/IPv6

- Source address
- Destination address
- Protocol
 - TCP/UDP/ICMP/ICMPv6
- TOS
- TTL (HLIMIT)

Supported on all kernels.

In OVS dump flows:

```
Ipv4:
ipv4(src=0.0.0.0/0.0.0.0, dst=0.0.0.0/0.0.0.0, proto=17, tos=0/0, ttl=0/0, frag=nc
Ipv6:
ipv6(src=::/::, dst=1:1:1::3:1040:1008, label=0/0, proto=58, tclass=0/0x3, h
```

Using TC rules:

```
IPv4:
tc filter add dev $rep parent ffff: protocol ip pref 1 \setminus
flower \
dst_ip 1.1.1.1 \
src_ip 1.1.1.2 \
ip_proto TCP \
ip_tos 0x3 ∖
ip_ttl 63 ∖
action mirred egress redirect dev $NIC
IPv6:
tc filter add dev $rep parent ffff: protocol ipv6 pref 1 \
flower \
dst_ip 1:1:1::3:1040:1009 \
src_ip 1:1:1::3:1040:1008 \
ip_proto TCP ∖
ip_tos 0x3 ∖
ip_ttl 63∖
action mirred egress redirect dev $NIC
```

TCP/UDP Source and Destination ports & TCP Flags

- TCP/UDP source and destinations ports
- TCP flags

Supported kernels are kernel > 4.13 and RHEL > 7.5

In OVS dump flows:

```
TCP: tcp(src=0/0,dst=32768/0x8000),
UDP: udp(src=0/0,dst=32768/0x8000),
TCP flags: tcp_flags(0/0)
```

Using TC rules:

```
tc filter add dev $rep parent ffff: protocol ip pref 1 \
flower \
ip_proto TCP \
dst_port 100 \
src_port 500 \
tcp_flags 0x4/0x7 \
action mirred egress redirect dev $NIC
```

VLAN

- ID
- Priority
- Inner vlan ID and Priority

Supported kernels: All (QinQ: kernel 4.19 and higher, and RHEL 7.7 and higher)

In OVS dump flows:

```
eth_type(0x8100), vlan(vid=2347, pcp=0),
```

Using TC rules:

```
tc filter add dev $rep parent ffff: protocol 802.1Q pref 1 \
                                          flower \
                                          vlan_ethtype 0x800 \
                                          vlan_id 100 \
                                          vlan_prio 0 \
                         action mirred egress redirect dev $NIC
QinQ:
tc filter add dev $rep parent ffff: protocol 802.1Q pref 1 \
                                          flower \
                                          vlan_ethtype 0x8100
\
                                          vlan_id 100 \
                                          vlan_prio 0 \
                                          cvlan_id 20 ∖
                                          cvlan_prio 0 \
                                          cvlan_ethtype 0x800 \
                         action mirred egress redirect dev $NIC
```

Tunnel

- ID (Key)
- Source IP address
- Destination IP address
- Destination port
- TOS (supported from kernel 4.19 and above & RHEL 7.7 and above)

- TTL (support from kernel 4.19 and above & RHEL 7.7 and above)
- Tunnel options (Geneve)

Supported kernels:

- VXLAN: All
- GRE: Kernel > 5.0, RHEL 7.7 and above
- Geneve: Kernel > 5.0, RHEL 7.7 and above

In OVS dump flows:

tunnel(tun_id=0x5, src=121.9.1.1, dst=131.10.1.1, ttl=0/0, tp_dst=4789, flags(+

Using TC rules:

```
# tc filter add dev $rep protocol 802.1Q parent ffff: pref 1
flower \
vlan_ethtype 0x800 \
vlan_id 100 \
vlan_prio 0 \
action mirred egress redirect dev $NIC
QinQ:
# tc filter add dev vxlan100 protocol ip parent ffff: \
                 flower \
                          skip_sw ∖
                          dst_mac e4:11:22:11:4a:51 \
                          src_mac e4+:11:22:11:4a:50 \
                          enc_src_ip 20.1.11.1 \
                          enc_dst_ip 20.1.12.1 \
                          enc_key_id 100 \
                          enc_dst_port 4789 \
                 action tunnel_key unset \
```

Supported Actions

Forward

Forward action allows for packet redirection:

- From VF to wire
- Wire to VF
- VF to VF

Supported on all kernels.

In OVS dump flows:

```
skb_priority(0/0), skb_mark(0/0), in_port(eth6), eth(src=00:02:10:40:10:{
, dst=68:54:ed:00:af:de), eth_type(0x8100), packets:1981, bytes:206024,
used:0.440s, dp:tc, actions:eth7
```

Using TC rules:

Drop

Drop action allows to drop incoming packets.

Supported on all kernels.

In OVS dump flows:

```
skb_priority(0/0), skb_mark(0/0), in_port(eth6), eth(src=00:02:10:40:10:{
,dst=68:54:ed:00:af:de), eth_type(0x8100), packets:1981, bytes:206024,
used:0.440s, dp:tc, actions:drop
```

Using TC rules:

tc filter add dev \$rep parent ffff: protocol arp pref 1	١
	flower $\$
	dst_mac
e4:1d:2d:5d:25:35 \	
	src_mac
e4:1d:2d:5d:25:34 \	
action drop	

Statistics

By default, each flow collects the following statistics:

- Packets number of packets which hit the flow
- Bytes total number of bytes which hit the flow
- Last used the amount of time passed since last packet hit the flow

Supported on all kernels.

```
skb_priority(0/0),skb_mark(0/0),in_port(eth6),eth(src=00:02:10:40:10:{
,dst=68:54:ed:00:af:de),eth_type(0x8100), packets:1981, bytes:206024,
used:0.440s, dp:tc, actions:drop
```

Using TC rules:

```
#tc -s filter show dev $rep ingress
filter protocol ip pref 2 flower chain 0
filter protocol ip pref 2 flower chain 0 handle 0x2
eth_type ipv4
ip_proto tcp
src_ip 192.168.140.100
src_port 80
skip_sw
in_hw
    action order 1: mirred (Egress Redirect to device p0v11_r)
stolen
    index 34 ref 1 bind 1 installed 144 sec used 0 sec
    Action statistics:
    Sent 388344 bytes 2942 pkt (dropped 0, overlimits 0 requeues 0)
    backlog 0b 0p requeues 0
```

Tunnels (Encapsulation/Decapsulation)

OVS-kernel supports offload of tunnels using encapsulation and decapsulation actions.

- Encapsulation pushing of tunnel header is supported on Tx
- Decapsulation popping of tunnel header is supported on Rx

Supported Tunnels:

- VXLAN (IPv4/IPv6) supported on all Kernels
- GRE (IPv4/IPv6) supported on kernel 5.0 and above & RHEL 7.6 and above
- Geneve (IPv4/IPv6) supported on kernel 5.0 and above & RHEL 7.6 and above

OVS configuration:

In case of offloading tunnel, the PF/bond should not be added as a port in the OVS datapath. It should rather be assigned with the IP address to be used for encapsulation.

The example below shows two hosts (PFs) with IPs 1.1.1.177 and 1.1.1.75, where the PF device on both hosts is enp4s0f0, and the VXLAN tunnel is set with VNID 98:

• On the first host:

```
# ip addr add 1.1.1.177/24 dev enp4s0f1
# ovs-vsctl add-port ovs-sriov vxlan0 -- set interface vxlan0
type=vxlan
options:local_ip=1.1.1.177 options:remote_ip=1.1.1.75
options:key=98
```

• On the second host:

```
# ip addr add 1.1.1.75/24 dev enp4s0f1
# ovs-vsctl add-port ovs-sriov vxlan0 -- set interface vxlan0
type=vxlan
options:local_ip=1.1.1.75 options:remote_ip=1.1.1.177
options:key=98
. for GRE IPv4 tunnel need use type=gre
. for GRE IPv6 tunnel need use type=ip6gre
```

j Note

When encapsulating guest traffic, the VF's device MTU must be reduced to allow the host/HW to add the encap headers without fragmenting the resulted packet. As such, the VF's MTU must be lowered by 50 bytes from the uplink MTU for IPv4 and 70 bytes for IPv6.

Tunnel offload using TC rules:

```
Encapsulation:
# tc filter add dev ens4f0_0 protocol 0x806 parent ffff: \
                 flower \
                         skip_sw ∖
                         dst_mac e4:11:22:11:4a:51 \
                         src_mac e4:11:22:11:4a:50 \
                 action tunnel_key set \
                 src_ip 20.1.12.1 \
                 dst_ip 20.1.11.1 \
                 id 100 \
                 action mirred egress redirect dev vxlan100
Decapsulation:
# tc filter add dev vxlan100 protocol 0x806 parent ffff: \
                 flower \
                           skip_sw ∖
                           dst_mac e4:11:22:11:4a:51 \
                           src_mac e4:11:22:11:4a:50 \
                           enc_src_ip 20.1.11.1 \
                           enc_dst_ip 20.1.12.1 \
```

enc_key_id 100 \
 enc_dst_port 4789 \
action tunnel_key unset \
 action mirred egress redirect dev ens4f0_0

VLAN Push/Pop

OVS-kernel supports offload of vlan header push/pop actions.

- Push—pushing of VLAN header is supported on Tx
- Pop—popping of tunnel header is supported on Rx

(i) Note

Starting with ConnectX-6 Dx hardware models and above, pushing of VLAN header is also supported on Rx, and popping of VLAN header is also supported on Tx.

OVS Configuration

Add a tag=\$TAG section for the OVS command line that adds the representor ports. For example, VLAN ID 52 is being used here.

```
# ovs-vsctl add-port ovs-sriov enp4s0f0
# ovs-vsctl add-port ovs-sriov enp4s0f0_0 tag=52
# ovs-vsctl add-port ovs-sriov enp4s0f0_1 tag=52
```

The PF port should not have a VLAN attached. This will cause OVS to add VLAN push/pop actions when managing traffic for these VFs.

Dump Flow Example

```
recirc_id(0),in_port(3),eth(src=e4:11:22:33:44:50,dst=00:02:c9:e9:bb:b
\
packets:0, bytes:0, used:never, actions:push_vlan(vid=52,pcp=0),2
recirc_id(0),in_port(2),eth(src=00:02:c9:e9:bb:b2,dst=e4:11:22:33:44:5
\
vlan(vid=52,pcp=0),encap(eth_type(0x0800),ipv4(frag=no)),
packets:0, bytes:0, used:never, actions:pop_vlan,3
```

VLAN Offload using TC Rules Example

```
# tc filter add dev ens4f0_0 protocol ip parent ffff: \
                 flower \
                         skip_sw ∖
                         dst_mac e4:11:22:11:4a:51 \
                         src_mac e4:11:22:11:4a:50 \
                 action vlan push id 100 \setminus
                 action mirred egress redirect dev ens4f0
# tc filter add dev ens4f0 protocol 802.1Q parent ffff: \
                 flower \
                         skip_sw ∖
                         dst_mac e4:11:22:11:4a:51 \
                         src_mac e4:11:22:11:4a:50 \
                         vlan_ethtype 0x800 \
                         vlan id 100 \
                         vlan_prio 0 \
                 action vlan pop \
                 action mirred egress redirect dev ens4f0_0
```

TC Configuration for ConnectX-6 Dx and Above

Example of VLAN Offloading with popping header on Tx and pushing on Rx using TC Rules:

```
# tc filter add dev ens4f0_0 ingress protocol 802.1Q parent ffff:
\
        flower \
                 vlan_id 100 \
         action vlan pop \setminus
         action tunnel_key set \
                 src_ip 4.4.4.1 \
                 dst_ip 4.4.4.2 \
                 dst_port 4789 ∖
                 id 42 \
        action mirred egress redirect dev vxlan0
# tc filter add dev vxlan0 ingress protocol all parent ffff: \
         flower \
                 enc_dst_ip 4.4.4.1 \
                 enc_src_ip 4.4.4.2 \
                 enc_dst_port 4789 \
                 enc_key_id 42 \setminus
        action tunnel_key unset \
         action vlan push id 100 \setminus
        action mirred egress redirect dev ens4f0_0
```

Header Rewrite

This action allows for modifying packet fields.

Ethernet Layer 2

- Destination MAC
- Source MAC

Supported kernels: Kernel 4.14 and above & RHEL 7.5 and above

```
skb_priority(0/0), skb_mark(0/0), in_port(eth6), eth(src=00:02:10:40:10:{
,dst=68:54:ed:00:af:de), eth_type(0x8100), packets:1981, bytes:206024,
used:0.440s, dp:tc, actions:
set(eth(src=68:54:ed:00:f4:ab,dst=fa:16:3e:dd:69:c4)), eth7
```

Using TC rules:

tc filter add dev \$rep parent ffff: pr	rotocol arp pref 1	flower $\$
e4:1d:2d:5d:25:35 \		dst_mac
e4.10.20.30.25.35 \		src_mac
e4:1d:2d:5d:25:34 \		
	action pedit ex	\
	munge eth dst se	et
20:22:33:44:55:66 \		
	munge eth src se	et
aa:ba:cc:dd:ee:fe \		
	action mirred eq	gress
redirect dev \$NIC		

IPv4/IPv6

- Source address
- Destination address
- Protocol
- TOS
- TTL (HLIMIT)

Supported kernels: Kernel 4.14 and above & RHEL 7.5 and above

In OVS dump flows:

```
Ipv4:
    set(eth(src=de:e8:ef:27:5e:45,dst=00:00:01:01:01:01)),
    set(ipv4(src=10.10.0.111,dst=10.20.0.122,ttl=63))
Ipv6:
    set(ipv6(dst=2001:1:6::92eb:fcbe:f1c8,hlimit=63)),
```

Using TC rules:

IPv4: tc filter add dev \$rep parent ffff: protocol ip pref 1	
	flower \ dst_ip
1.1.1.1 \	src_ip
1.1.1.2 \ TCP \	ip_proto
	ip_tos
	ip_ttl 63
` pedit ex ∖	
munge ip src se	et 2.2.2.1 \
munge ip dst se	et 2.2.2.2 \
munge ip tos se	et 0 \
munge ip ttl de	ec /
action mirred e	egress
redirect dev \$NIC	

IPv6: tc filter add dev \$rep parent ffff: protocol ipv6 pref 1 \setminus flower \ dst_ip 1:1:1::3:1040:1009 \ src_ip 1:1:1::3:1040:1008 \ ip_proto tcp ∖ ip_tos 0x3 \ ip_ttl 63\ pedit ex \ munge ipv6 src set 2:2:2:3:1040:1009 \ munge ipv6 dst set 2:2:2:3:1040:1008 \ munge ipv6 hlimit dec $\$ action mirred egress redirect dev \$NIC

i Note

IPv4 and IPv6 header rewrite is only supported with match on UDP/TCP/ICMP protocols.

TCP/UDP Source and Destination Ports

• TCP/UDP source and destinations ports

Supported kernels: kernel > 4.16 & RHEL > 7.6

In OVS dump flows:

TCP:	<pre>set(tcp(src=</pre>	32768/0xffff,dst=32768/0xffff)),
UDP:		
	<pre>set(udp(src=</pre>	32768/0xffff,dst=32768/0xffff)),

Using TC rules:

TCP:						
pref 1 \	tc filter	add dev	\$rep	parent ff	f: pro	tocol ip
						flower \ dst_ip
1.1.1.1 \						
						src_ip
1.1.1.2 \						
tcp \						ip_proto
						ip_tos
0x3 \						- -
						ip_ttl 63
\						
				pedit ex		
				pedit ex	munge	ip tcp
sport set 200						in ton
dport set 200				pedit ex	munge	тр гср
aport Set 200				action m	irred e	aress
redirect dev \$N	IC			2002011		<u> </u>

UDP:	tc filter add dev \$rep parent ffff: pr	otocol in
pref 1 \		010001 10
		flower \
1.1.1.1 \		dst_ip
		src_ip
1.1 . 1.2 \		ip_proto
udp \		
0x3 \		ip_tos
		ip_ttl 63
\		
	pedit ex \ pedit ex munge	ip udp
sport set 200	poure on mange	<u>-</u> p aap
	pedit ex munge	ip udp
dport set 200	action mirred	earess
redirect dev \$N		

VLAN

• ID

Supported on all kernels.

In OVS dump flows:

Set(vlan(vid=2347,pcp=0/0)),

Using TC rules:

Connection Tracking

The TC connection tracking action performs connection tracking lookup by sending the packet to netfilter conntrack module. Newly added connections may be associated, via the ct commit action, with a 32 bit mark, 128 bit label and source/destination NAT values.

The following example allows ingress tcp traffic from the uplink representor to vf1_rep, while assuring that egress traffic from vf1_rep is only allowed on established connections. In addition, mark and source IP NAT is applied.

In OVS dump flows:

```
ct(zone=2, nat)
ct_state(+est+trk)
actions:ct(commit,zone=2,mark=0x4/0xfffffff,nat(src=5.5.5.5))
```

Using TC rules:

```
# tc filter add dev $uplink_rep ingress chain 0 prio 1 proto ip \
        flower \
        ip_proto tcp \
        ct_state -trk \
        action ct zone 2 nat pipe
        action goto chain 2
```

```
# tc filter add dev $uplink_rep ingress chain 2 prio 1 proto ip
\
                      flower \
                      ct_state +trk+new \
                action ct zone 2 commit mark Oxbb nat src addr
5.5.5.7 pipe \
                action mirred egress redirect dev $vf1_rep
# tc filter add dev $uplink_rep ingress chain 2 prio 1 proto ip
                     flower \
                     ct_zone 2 \setminus
                     ct_mark Oxbb \
                     ct_state +trk+est \
                 action mirred egress redirect dev $vf1_rep
  #Setup filters on $vf1_rep, allowing only established
connections of zone 2 through, and reverse nat (dst nat in
this case)
# tc filter add dev $vf1_rep ingress chain 0 prio 1 proto ip \
                      flower \
                      ip_proto tcp \
                      ct_state -trk \
                 action ct zone 2 nat pipe \setminus
                 action goto chain 1
  # tc filter add dev $vf1_rep ingress chain 1 prio 1 proto ip \
                      flower \
                      ct_zone 2 \setminus
                      ct_mark Oxbb \
                      ct_state +trk+est \
                  action mirred egress redirect dev eth0
```

Connection Tracking Performance Tuning

• Max offloaded connections—specifies the limit on the number of offloaded connections.

```
# devlink dev param set pci/${pci_dev} name
ct_max_offloaded_conns value $max cmode runtime
```

• Allow mixed NAT/non-NAT CT—allows offloading of the following scenario:

```
cookie=0x0, duration=21.843s, table=0,
n_packets=4838718, n_bytes=241958846, ct_state=-
trk,ip,in_port=enp8s0f0 actions=ct(table=1,zone=2)
cookie=0x0, duration=21.823s, table=1, n_packets=15363,
n_bytes=773526, ct_state=+new+trk,ip,in_port=enp8s0f0
actions=ct(commit,zone=2,nat(dst=11.11.11.11)),output:"enp8s0f0_1"
cookie=0x0, duration=21.806s, table=1,
n_packets=4767594, n_bytes=238401190,
ct_state=+est+trk,ip,in_port=enp8s0f0
actions=ct(zone=2,nat),output:"enp8s0f0_1"
```

Example:

```
# echo enable >
/sys/class/net/<device>/compat/devlink/ct_action_on_nat_conns
```

Forward to Chain (TC Only)

TC interface supports adding flows on different chains. Only chain 0 is accessed by default. Access to the other chains requires usafe of the goto action.

In this example, a flow is created on chain 1 without any match and redirect to wire.

The second flow is created on chain 0 and match on source MAC and action goto chain 1.

This example simulates simple MAC spoofing.

Port Mirroring (Flow Based VF Traffic Mirroring for ASAP²)

Unlike para-virtual configurations, when the VM traffic is offloaded to the hardware via SR-IOV VF, the host side Admin cannot snoop the traffic (e.g. for monitoring).

ASAP² uses the existing mirroring support in OVS and TC along with the enhancement to the offloading logic in the driver to allow mirroring the VF traffic to another VF.

The mirrored VF can be used to run traffic analyzer (tcpdump, wireshark, etc) and observe the traffic of the VF being mirrored.

The example below shows the creation of port mirror on the following configuration:

```
# ovs-vsctl show
09d8a574-9c39-465c-9f16-47d81c12f88a
Bridge br-vxlan
Port "enp4s0f0_1"
Interface "enp4s0f0_1"
```

```
Port "vxlan0"
Interface "vxlan0"
type: vxlan
options: {key="100",
remote_ip="192.168.1.14"}
Port "enp4s0f0_0"
Interface "enp4s0f0_0"
Port "enp4s0f0_2"
Interface "enp4s0f0_2"
Port br-vxlan
Interface br-vxlan
type: internal
ovs_version: "2.14.1"
```

• To set **enp4s0f0_0** as the mirror port, and mirror all of the traffic:

• To set **enp4sOfO_O** as the mirror port, and only mirror the traffic, the destination is enp4sOfO_1:

• To set enp4s0f0_0 as the mirror port, and only mirror the traffic the source is enp4s0f0_1:

• To set enp4s0f0_0 as the mirror port and mirror, all the traffic on enp4s0f0_1:

To clear the mirror port:

ovs-vsctl clear bridge br-vxlan mirrors

Mirroring using TC:

```
action mirred
egress redirect dev $NIC
Mirror to tunnel:
tc filter add dev $rep parent ffff: protocol arp pref 1 \setminus
                                                    flower \
                                                    dst_mac
e4:1d:2d:5d:25:35 \
                                                    src_mac
e4:1d:2d:5d:25:34 \
                                  action tunnel_key set \
                                  src_ip 1.1.1.1 \
                                  dst_ip 1.1.1.2 \
                                  dst_port 4789 \
                                  id 768 \
                                  pipe \
                                  action mirred egress mirror dev
vxlan100 pipe \
                                  action mirred egress redirect dev
$NIC
```

Forward to Multiple Destinations

Forward to up 32 destinations (representors and tunnels) is supported using TC.

Example 1: forward to 32 VFs.

```
tc filter add dev $NIC parent ffff: protocol arp pref 1 \
        flower \
        dst_mac e4:1d:2d:5d:25:35 \
        src_mac e4:1d:2d:5d:25:34 \
        action mirred egress mirror dev $rep0
pipe \
```

	action	mirred	egress	mirror	dev	\$rep1
pipe \						
	action	mirred	egress	mirror	dev	\$rep30
pipe \						
	action	mirred	egress	redired	ct de	ev \$rep31

Example 2: forward to 16 tunnels.

```
tc filter add dev $rep parent ffff: protocol arp pref 1 \
                        flower \
                        dst_mac e4:1d:2d:5d:25:35 \
                        src_mac e4:1d:2d:5d:25:34 \
                        action tunnel_key set src_ip $ip_src
dst_ip $ip_dst \
dst_port 4789 id 0 nocsum \
pipe action mirred egress mirror dev vxlan0 pipe \
                        action tunnel_key set src_ip $ip_src
dst_ip $ip_dst \
dst_port 4789 id 1 nocsum \
pipe action mirred egress mirror dev vxlan0 pipe \
. . .
                        action tunnel_key set src_ip $ip_src
dst_ip $ip_dst \
dst_port 4789 id 15 nocsum \
pipe action mirred egress redirect dev vxlan0
```

(i) Note

- TC supports up to 32 actions
- If header rewrite is used, then all destinations should have the same header rewrite

• If VLAN push/pop is used, then all destinations should have the same VLAN ID and actions

sFLOW

This feature allows for monitoring traffic sent between two VMs on the same host using an sFlow collector.

The example below assumes the environment is configured as described below.

```
# ovs-vsctl show
  09d8a574-9c39-465c-9f16-47d81c12f88a
       Bridge br-vxlan
                                       Port "enp4s0f0_1"
                                              Interface "enp4s0f0_1"
                                       Port "vxlan0"
                                              Interface "vxlan0"
type: vxlan
options: {key="100", remote_ip="192.168.1.14"}
                                       Port "enp4s0f0_0"
                                              Interface "enp4s0f0_0"
                                       Port "enp4s0f0_2"
                                              Interface "enp4s0f0_2"
                                       Port br-vxlan
                                              Interface br-vxlan
type: internal
      ovs_version: "2.14.1"
```

To sample all traffic over the OVS bridge:

```
# ovs-vsctl -- --id=@sflow create sflow agent=\"$SFLOW_AGENT\" \
        target=\"$SFLOW_TARGET:$SFLOW_PORT\"
header=$SFLOW_HEADER \
        sampling=$SFLOW_SAMPLING polling=10 \
        -- set bridge br-vxlan sflow=@sflow
```

Parameter	Description
SFLOW_AGENT	Indicates that the sFlow agent should send traffic from SFLOW_AGENT's IP address
SFLOW_TARGE T	Remote IP address of the sFLOW collector
SFLOW_HEADE R	Size of packet header to sample (in bytes)
SFLOW_SAMPL ING	Sample rate

To clear the sFLOW configuration:

```
# ovs-vsctl clear bridge br-vxlan sflow
```

To list the sFLOW configuration:

ovs-vsctl list sflow

sFLOW using TC:

```
Sample to VF
tc filter add dev $rep parent ffff: protocol arp pref 1 \
flower \
```

e4:1d:2d:5d:25:35 \
e4:1d:2d:5d:25:34 \
rate 10 group 5 trunc 96 \
egress redirect dev \$NIC

dst_mac

src_mac

action sample

action mirred

(i) Note

Userspace application is needed in order to process to sampled packet from the kernel. Example: <u>https://github.com/Mellanox/libpsample</u>

Rate Limit

OVS-kernel supports offload of VF rate limit using OVS configuration and TC.

The example below sets a rate limit to the VF related to representor eth0 to 10Mbps.

```
OVS:
  # ovs-vsctl set interface eth0 ingress_policing_rate=10000
tc:
  # tc_filter add dev eth0 root prio 1 protocol ip matchall
skip_sw action police rate 10mbit burst 20k
```

Kernel Requirements

This kernel config should be enabled in order to support switchdev offload.

- CONFIG_NET_ACT_CSUM needed for action csum
- CONFIG_NET_ACT_PEDIT needed for header rewrite
- CONFIG_NET_ACT_MIRRED needed for basic forward
- CONFIG_NET_ACT_CT needed for connection tracking (supported from kernel 5.6)
- CONFIG_NET_ACT_VLAN needed for action vlan push/pop
- CONFIG_NET_ACT_GACT
- CONFIG_NET_CLS_FLOWER
- CONFIG_NET_CLS_ACT
- CONFIG_NET_SWITCHDEV
- CONFIG_NET_TC_SKB_EXT needed for connection tracking (supported from kernel 5.6)
- CONFIG_NET_ACT_CT needed for connection tracking (supported from kernel 5.6)
- CONFIG_NFT_FLOW_OFFLOAD
- CONFIG_NET_ACT_TUNNEL_KEY
- CONFIG_NF_FLOW_TABLE needed for connection tracking (supported from kernel 5.6)
- CONFIG_SKB_EXTENSIONS needed for connection tracking (supported from kernel 5.6)
- CONFIG_NET_CLS_MATCHALL
- CONFIG_NET_ACT_POLICE
- CONFIG_MLX5_ESWITCH

VF Metering

OVS-kernel supports offloading of VF metering (TX and RX) using sysfs. Metering of number of packets per second (PPS) and bytes per second (BPS) is supported.

The example bellow sets Rx meter on VF 0 with value 10Mbps BPS.

```
echo 10000000 >
/sys/class/net/enp4s0f0/device/sriov/0/meters/rx/bps/rate
echo 65536 >
/sys/class/net/enp4s0f0/device/sriov/0/meters/rx/bps/burst
```

The example bellow sets Tx meter on VF 0 with value 1000 PPS.

```
echo 1000 >
/sys/class/net/enp4s0f0/device/sriov/0/meters/tx/pps/rate
echo 100 >
/sys/class/net/enp4s0f0/device/sriov/0/meters/tx/pps/burst
```

i Note

Both rate and burst must not be zero and burst may need to be adjusted according to the requirements.

The following counters can be used to query the number dropped packet/bytes:

#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/rx/pps/packets_dropped
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/rx/pps/bytes_dropped
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/rx/bps/packets_dropped

```
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/rx/bps/bytes_dropped
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/tx/pps/bytes_dropped
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/tx/bps/packets_dropped
#cat
/sys/class/net/enp8s0f0/device/sriov/0/meters/tx/bps/packets_dropped
```

Representor Metering

Metering for uplink and VF representors traffic support has been added.

Traffic going to a representor device can be a result of a miss in the embedded switch (eSwitch) FDB tables. This means that a packet which arrived from that representor into the eSwitch was not matched against the existing rules in the hardware FDB tables and needs to be forwarded to software to be handled there and is, therefore, forwarded to the originating representor device driver.

The meter allows to configure the max rate [packets/sec] and max burst [packets] for traffic going to the representor driver. Any traffic exceeding values provided by the user will be dropped in hardware. There are statistics that show number of dropped packets.

The configuration of a representors metering is done via a new sysfs called miss_rl_cfg.

- Full path of the miss_rl_cfg parameter: /sys/class/net//rep_config/miss_rl_cfg
- Usage: echo "<rate> <burst>" > /sys/class/net//rep_config/miss_rl_cfg. Rate is the max rate of packets allowed for this representor (in packets/sec units) and burst is the max burst size allowed for this representor (in packets units). Both values must be specified. The default is 0 for both, meaning unlimited rate and burst.

To view the amount of packets and bytes that were dropped due to traffic exceeding the user-provided rate and burst, two read-only sysfs for statistics are exposed.

• /sys/class/net//rep_config/miss_rl_dropped_bytes counts how many FDB-miss bytes were dropped due to reaching the miss limits

• /sys/class/net//rep_config/miss_rl_dropped_packets counts how many FDB-miss packets were dropped due to reaching the miss limits

Open vSwitch Metering

There are two types of meters, kpps (kilobits per second) and pktps(packets per second), which are described in Meter Syntax of OpenFlow 1.3+ Switch Meter Table Commands. OVS-Kernel supports offloading them both.

The example below is to offload a kpps meter. Please follow the steps after doing basic configurations as described in section 5.1.3.

1. Create OVS meter with a target rate.

```
ovs-ofctl -0 OpenFlow13 add-meter ovs-sriov
meter=1,kbps,band=type=drop,rate=204800
```

2. Delete the default rule.

```
ovs-ofctl del-flows ovs-sriov
```

3. Configure OpenFlow rules. Here VF bandwidth on the receiving side will be limited by the rate configured in step 1.

```
ovs-ofctl -0 OpenFlow13 add-flow ovs-sriov
'ip,dl_dst=e4:11:22:33:44:50,actions= meter:1,output:enp4s0f0_0'
    ovs-ofctl -0 OpenFlow13 add-flow ovs-sriov
'ip,dl_src=e4:11:22:33:44:50,actions= output:enp4s0f0'
    ovs-ofctl -0 OpenFlow13 add-flow ovs-sriov 'arp,actions=normal'
```

4. Run iperf server and be ready to receive UDP traffic. On the outer node, run iperf client to send UDP traffic to this VF. After traffic starts, check the offloaded meter rule.

```
ovs-appctl dpctl/dump-flows --names type=offloaded
recirc_id(0),in_port(enp4s0f0),eth(dst=e4:11:22:33:44:50),eth_typ
packets:11626587, bytes:17625889188, used:0.470s,
actions:meter(0),enp4s0f0_0
```

In order to verify metering, iperf client should set the target bandwidth with a number which is larger than the meter rate configured. Then it will be visible that packets are received with the limited rate on the server side and the extra packets are dropped by hardware.

Multiport eSwitch Mode

The multiport eSwitch mode allows to add rules on a VF representor with an action forwarding the packet to the physical port of the physical function. This can be used to implement failover or forward packets based on external information such the cost of the route.

- 1. To configure this more, the nvconig parameter LAG_RESOURCE_ALLOCATION must be set.
- 2. After the driver loads, configure multiport eSwitch for each PF where enp8s0f0 and enp8s0f1 represent the netdevices for the PFs.

```
echo multiport_esw >
/sys/class/net/enp8s0f0/compat/devlink/lag_port_select_mode
echo multiport_esw >
/sys/class/net/enp8s0f1/compat/devlink/lag_port_select_mode
```

The mode becomes operational after entering switchdev mode on both PFs.

Rule example:

tc filter add dev enp8s0f0_0 prot ip root flower dst_ip 7.7.7.7 action mirred egress redirect dev enp8s0f1

OVS-DPDK Hardware Offloads

OVS-DPDK Hardware Offloads Configuration

To configure OVS-DPDK HW offloads:

1. Unbind the VFs.

echo 0000:04:00.2 > /sys/bus/pci/drivers/mlx5_core/unbind echo 0000:04:00.3 > /sys/bus/pci/drivers/mlx5_core/unbind

Note: VMs with attached VFs must be powered off to be able to unbind the VFs.

2. Change the e-switch mode from Legacy to SwitchDev on the PF device (make sure all VFs are unbound). This will also create the VF representor netdevices in the host OS.

echo switchdev > /sys/class/net/enp4s0f0/compat/devlink/mode

To revert to SR-IOV Legacy mode:

echo legacy > /sys/class/net/enp4s0f0/compat/devlink/mode

Note that running this command will also result in the removal of the VF representor netdevices.

3. Bind the VFs.

echo 0000:04:00.2 > /sys/bus/pci/drivers/mlx5_core/bind echo 0000:04:00.3 > /sys/bus/pci/drivers/mlx5_core/bind

4. Run the Open vSwitch service.

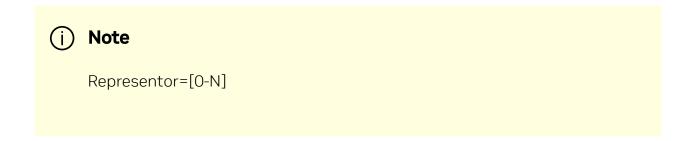
systemctl start openvswitch

5. Enable hardware offload (disabled by default).

```
ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk-
init=true
ovs-vsctl set Open_vSwitch . other_config:hw-offload=true
```

6. Configure the DPDK white list.

```
ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk-
extra="-a 0000:01:00.0,representor=[0],dv_flow_en=1,dv_esw_en=1,dv_xmeta_en=1"
```



7. Restart the Open vSwitch service. This step is required for HW offload changes to take effect.

```
systemctl restart openvswitch
```

8. Create OVS-DPDK bridge.

```
ovs-vsctl --no-wait add-br br0-ovs -- set bridge br0-ovs
datapath_type=netdev
```

9. Add PF to OVS.

```
ovs-vsctl add-port br0-ovs pf -- set Interface pf type=dpdk
options:dpdk-devargs=0000:88:00.0
```

10. Add representor to OVS.

```
ovs-vsctl add-port br0-ovs representor -- set Interface
representor type=dpdk options:dpdk-
devargs=0000:88:00.0, representor=[0]
```



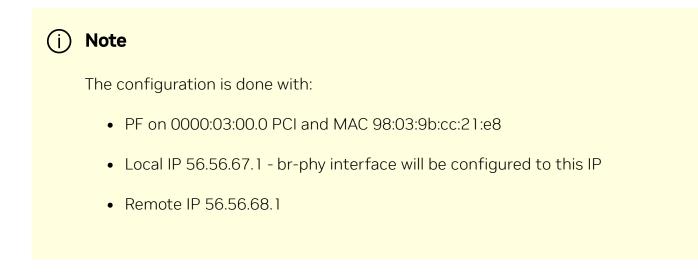
Offloading VXLAN Encapsulation/Decapsulation Actions

vSwitch in userspace rather than kernel-based Open vSwitch requires an additional bridge. The purpose of this bridge is to allow use of the kernel network stack for routing

and ARP resolution.

The datapath needs to look-up the routing table and ARP table to prepare the tunnel header and transmit data to the output port.

Configuring VXLAN Encap/Decap Offloads



To configure OVS-DPDK VXLAN:

1. Create a br-phy bridge.

```
ovs-vsctl add-br br-phy -- set Bridge br-phy
datapath_type=netdev -- br-set-external-id br-phy bridge-id
br-phy -- set bridge br-phy fail-mode=standalone
other_config:hwaddr=98:03:9b:cc:21:e8
```

2. Attach PF interface to br-phy bridge.

ovs-vsctl add-port br-phy p0 -- set Interface p0 type=dpdk
options:dpdk-devargs=0000:03:00.0

3. Configure IP to the bridge.

ip addr add 56.56.67.1/24 dev br-phy

4. Create a br-ovs bridge.

```
ovs-vsctl add-br br-ovs -- set Bridge br-ovs
datapath_type=netdev -- br-set-external-id br-ovs bridge-id
br-ovs -- set bridge br-ovs fail-mode=standalone
```

5. Attach representor to br-ovs.

```
ovs-vsctl add-port br-ovs pf0vf0 -- set Interface pf0vf0
type=dpdk options:dpdk-devargs=0000:03:00.0,representor=[0]
```

6. Add a port for the VXLAN tunnel.

```
ovs-vsctl add-port ovs-sriov vxlan0 -- set interface vxlan0
type=vxlan options:local_ip=56.56.67.1
options:remote_ip=56.56.68.1 options:key=45 options:dst_port=4789
```

Connection Tracking Offload

Connection tracking enables stateful packet processing by keeping a record of currently open connections.

OVS flows using connection tracking can be accelerated using advanced Network Interface Cards (NICs) by offloading established connections.

To view offloaded connections, run:

```
ovs-appctl dpctl/offload-stats-show
```

SR-IOV VF LAG

To configure OVS-DPDK SR-IOV VF LAG:

1. Enable SR-IOV on the NICs.

mlxconfig -d <PCI> set SRIOV_EN=1

2. Allocate the desired number of VFs per port.

```
echo $n > /sys/class/net/<net name>/device/sriov_numvfs
```

3. Unbind all VFs.

```
echo <VF PCI> >/sys/bus/pci/drivers/mlx5_core/unbind
```

4. Change both NICs' mode to SwitchDev.

devlink dev eswitch set pci/<PCI> mode switchdev

5. Create Linux bonding using kernel modules.

modprobe bonding mode=<desired mode>

Note: Other bonding parameters can be added here. The supported Bond modes are: Active-Backup, XOR and LACP.

6. Bring all PFs and VFs down.

```
ip link set <PF/VF> down
```

7. Attach both PFs to the bond.

ip link set <PF> master bond0

8. To work with VF-LAG with OVS-DPDK, add the bond master (PF) to the bridge.

```
ovs-vsctl add-port br-phy p0 -- set Interface p0 type=dpdk
options:dpdk-devargs=0000:03:00.0 options:dpdk-lsc-
interrupt=true
```

9. Add representor \$N of PFO or PF1 to a bridge.

ovs-vsctl add-port br-phy rep\$N -- set Interface rep\$N
type=dpdk options:dpdk-devargs=<PF0 PCI>,representor=pf0vf\$N

OR

```
ovs-vsctl add-port br-phy rep$N -- set Interface rep$N
type=dpdk options:dpdk-devargs=<PF0 PCI>,representor=pf1vf$N
```

VirtIO Acceleration through VF Relay (Software & Hardware vDPA)

(i) Note

Hardware vDPA is supported on ConnectX-6 Dx, ConnectX-6 Lx & BlueField-2 cards and above only.

i) Note

Hardware vDPA is enabled by default. In case your hardware does not support vDPA, the driver will fall back to Software vDPA.

To check which vDPA mode is activated on your driver, run: ovs-ofctl -O OpenFlow14 dump-ports br0-ovs and look for hw-mode flag.

(i) Note

This feature has not been accepted to the OVS-DPDK Upstream yet, making its API subject to change.

In user space, there are two main approaches for communicating with a guest (VM), either through SR-IOV, or through virtIO.

Phy ports (SR-IOV) allow working with port representor, which is attached to the OVS and a matching VF is given with pass-through to the guest. HW rules can process packets from up-link and direct them to the VF without going through SW (OVS). Therefore, using SR-IOV achieves the best performance.

However, SR-IOV architecture requires the guest to use a driver specific to the underlying HW. Specific HW driver has two main drawbacks:

- 1. Breaks virtualization in some sense (guest is aware of the HW). It can also limit the type of images supported.
- 2. Gives less natural support for live migration.

Using virtIO port solves both problems. However, it reduces performance and causes loss of some functionalities, such as, for some HW offloads, working directly with virtIO. To solve this conflict, a new netdev type- dpdkvdpa has been created. The new netdev is similar to the regular DPDK netdev, yet introduces several additional functionalities.

dpdkvdpa translates between phy port to virtIO port. It takes packets from the Rx queue and sends them to the suitable Tx queue, and allows transfer of packets from virtIO guest (VM) to a VF, and vice-versa, benefitting from both SR-IOV and virtIO.

To add vDPA port:

```
ovs-vsctl add-port br0 vdpa0 -- set Interface vdpa0 type=dpdkvdpa
\
options:vdpa-socket-path=<sock path> \
options:vdpa-accelerator-devargs=<vf pci id> \
options:dpdk-devargs=<pf pci id>,representor=[id] \
options: vdpa-max-queues =<num queues> \
options: vdpa-sw=<true/false>
```

Note: vdpa-max-queues is an optional field. When the user wants to configure 32 vDPA ports, the maximum queues number is limited to 8.

vDPA Configuration in OVS-DPDK Mode

Prior to configuring vDPA in OVS-DPDK mode, follow the steps below.

1. Generate the VF.

echo 0 > /sys/class/net/enp175s0f0/device/sriov_numvfs
echo 4 > /sys/class/net/enp175s0f0/device/sriov_numvfs

2. Unbind each VF.

```
echo <pci> > /sys/bus/pci/drivers/mlx5_core/unbind
```

3. Switch to SwitchDev mode.

echo switchdev >> /sys/class/net/enp175s0f0/compat/devlink/mode

4. Bind each VF.

```
echo <pci> > /sys/bus/pci/drivers/mlx5_core/bind
```

5. Initialize OVS with:

```
ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk-
init=true
ovs-vsctl --no-wait set Open_vSwitch . other_config:hw-
offload=true
```

To configure vDPA in OVS-DPDK mode on ConnectX-5 cards and above:

1. Open vSwitch configuration.

```
ovs-vsctl --no-wait set Open_vSwitch . other_config:dpdk-
extra="-a 0000:01:00.0,representor=[0],dv_flow_en=1,dv_esw_en=1,dv_xmeta_en=1"
/usr/share/openvswitch/scripts/ovs-ctl restart
```

2. Create OVS-DPDK bridge.

```
ovs-vsctl add-br br0-ovs -- set bridge br0-ovs
datapath_type=netdev
```

```
ovs-vsctl add-port br0-ovs pf -- set Interface pf type=dpdk
options:dpdk-devargs=0000:01:00.0
```

3. Create vDPA port as part of the OVS-DPDK bridge.

```
ovs-vsctl add-port br0-ovs vdpa0 -- set Interface vdpa0
type=dpdkvdpa options:vdpa-socket-path=/var/run/virtio-
forwarder/sock0 options:vdpa-accelerator-devargs=0000:01:00.2
options:dpdk-devargs=0000:01:00.0,representor=[0] options:
vdpa-max-queues=8
```

To configure vDPA in OVS-DPDK mode on BlueField cards:

Set the bridge with the software or hardware vDPA port:

• On the ARM side:

Create the OVS-DPDK bridge.

```
ovs-vsctl add-br br0-ovs -- set bridge br0-ovs
datapath_type=netdev
ovs-vsctl add-port br0-ovs pf -- set Interface pf type=dpdk
options:dpdk-devargs=0000:af:00.0
ovs-vsctl add-port br0-ovs rep-- set Interface rep type=dpdk
options:dpdk-devargs=0000:af:00.0,representor=[0]
```

• On the host side:

Create the OVS-DPDK bridge.

```
ovs-vsctl add-br br1-ovs -- set bridge br1-ovs
datapath_type=netdev protocols=0penFlow14
```

ovs-vsctl add-port br0-ovs vdpa0 -- set Interface vdpa0
type=dpdkvdpa options:vdpa-socket-path=/var/run/virtioforwarder/sock0 options:vdpa-accelerator-devargs=0000:af:00.2

Note: To configure **SW** vDPA, add "options:vdpa-sw=true" to the end of the command.

Software vDPA Configuration in OVS-Kernel Mode

SW vDPA can also be used in configurations where the HW offload is done through TC and not DPDK.

1. Open vSwitch configuration.

```
ovs-vsctl set Open_vSwitch . other_config:dpdk-extra="-a
0000:01:00.0,representor=[0],dv_flow_en=1,dv_esw_en=0,idv_xmeta_en=0,isolated_mode=1"
/usr/share/openvswitch/scripts/ovs-ctl restart
```

2. Create OVS-DPDK bridge.

```
ovs-vsctl add-br br0-ovs -- set bridge br0-ovs
datapath_type=netdev
```

3. Create vDPA port as part of the OVS-DPDK bridge.

```
ovs-vsctl add-port br0-ovs vdpa0 -- set Interface vdpa0
type=dpdkvdpa options:vdpa-socket-path=/var/run/virtio-
forwarder/sock0 options:vdpa-accelerator-devargs=0000:01:00.2
options:dpdk-devargs=0000:01:00.0,representor=[0] options:
vdpa-max-queues=8
```

4. Create Kernel bridge.

```
ovs-vsctl add-br br-kernel
```

5. Add representors to Kernel bridge.

```
ovs-vsctl add-port br-kernel enp1s0f0_0
ovs-vsctl add-port br-kernel enp1s0f0
```

Large MTU/Jumbo Frame Configuration

To configure MTU/jumbo frames:

1. Verify that the Kernel version on the VM is 4.14 or above.

```
cat /etc/redhat-release
```

2. Set the MTU on both physical interfaces in the host.

ifconfig ens4f0 mtu 9216

3. Send a large size packet and verify that it is sent and received correctly.

```
tcpdump -i ens4f0 -nev icmp & ping 11.100.126.1 -s 9188 -M do -c 1
```

4. Enable host_mtu in xml, and add the following values to xml.

```
host_mtu=9216, csum=on, guest_csum=on, host_tso4=on, host_tso6=on
```

Example:

```
<qemu:commandline>
<qemu:arg value='-chardev'/>
<qemu:arg value='socket,id=charnet1,path=/tmp/sock0,server'/>
<qemu:arg value='-netdev'/>
<qemu:arg value='vhost-user,chardev=charnet1,queues=16,id=hostnet1'/>
<qemu:arg value='user'/>
<qemu:arg value='user'/
<qemu:arg value='user'/>
<qemu:arg value='user'/
<qemu:arg value='user'/
<qemu:arg value='user'/
<qemu:arg value='user'/
<qemu:arg value='user'/
<qem
```

5. Add mtu_request=9216 option to the OvS ports inside the container and restart the OVS:

ovs-vsctl add-port br0-ovs pf -- set Interface pf type=dpdk
options:dpdk-devargs=0000:c4:00.0 mtu_request=9216

OR:

ovs-vsctl add-port br0-ovs vdpa0 -- set Interface vdpa0 type=dpdkvdpa options:vdpa-socket-path=/tmp/sock0 options:vdpa-accelerator-devargs=0000:c4:00.2 options:dpdkdevargs=0000:c4:00.0, representor=[0] mtu_request=9216 /usr/share/openvswitch/scripts/ovs-ctl restart

6. Start the VM and configure the MTU on the VM.

ifconfig eth0 11.100.124.2/16 up ifconfig eth0 mtu 9216 ping 11.100.126.1 -s 9188 -M do -c1

E2E Cache

i) Note

This feature is at beta level.

OVS offload rules are based on a multi-table architecture. E2E cache feature enables merging the multi-table flow matches and actions into one joint flow.

This improves connection tracking performance by using a single-table when exact match is detected.

• To set the E2E cache size (default = 4k):

ovs-vsctl set open_vswitch . other_config:e2e-size=<size>
systemctl restart openvswitch

Note: Make sure to restart the openvswitch service in order for the configuration to take effect.

• To enable/disable E2E cache (default = disabled) :

```
ovs-vsctl set open_vswitch . other_config:e2e-enable=
<true/false>
systemctl restart openvswitch
```

Note: Make sure to restart the openvswitch service in order for the configuration to take effect.

• To run E2E cache statistics:

ovs-appctl dpctl/dump-e2e-stats

• To run E2E cache flows:

ovs-appctl dpctl/dump-e2e-flows

Geneve Encapsulation/Decapsulation

Geneve tunneling offload feature support includes matching on extension header.

To configure OVS-DPDK Geneve encap/decap:

1. Create a br-phy bridge.

```
ovs-vsctl --may-exist add-br br-phy -- set Bridge br-phy
datapath_type=netdev -- br-set-external-id br-phy bridge-id
br-phy -- set bridge br-phy fail-mode=standalone
```

2. Attach PF interface to br-phy bridge.

```
ovs-vsctl add-port br-phy pf -- set Interface pf type=dpdk
options:dpdk-devargs=<PF PCI>
```

3. Configure IP to the bridge.

```
ifconfig br-phy <$local_ip_1> up
```

4. Create a br-int bridge.

```
ovs-vsctl --may-exist add-br br-int -- set Bridge br-int
datapath_type=netdev -- br-set-external-id br-int bridge-id
br-int -- set bridge br-int fail-mode=standalone
```

5. Attach representor to br-int.

```
ovs-vsctl add-port br-int rep$x -- set Interface rep$x
type=dpdk options:dpdk-devargs=<PF PCI>,representor=[$x]
```

6. Add a port for the GENEVE tunnel.

```
ovs-vsctl add-port br-int geneve0 -- set interface geneve0
type=geneve options:key=<VNI> options:remote_ip=
<$remote_ip_1> options:local_ip=<$local_ip_1>
```

Parallel Offloads

OVS-DPDK supports parallel insertion and deletion of offloads (flow & CT). While multiple threads are supported, by default only one is used.

To configure multiple threads:

```
ovs-vsctl set Open_vSwitch . other_config:n-offload-threads=3
```

Make sure to restart the openvswitch service in order for the configuration to take effect.

systemctl restart openvswitch

i) Note

For more information, see the OvS user manual.

sFlow

This feature allows for monitoring traffic sent between two VMs on the same host using an sFlow collector.

To sample all traffic over the OVS bridge, run the following:

```
# ovs-vsctl -- --id=@sflow create sflow agent=\"$SFLOW_AGENT\" \
        target=\"$SFLOW_TARGET:$SFLOW_HEADER\"
header=$SFLOW_HEADER \
        sampling=$SFLOW_SAMPLING polling=10 \
        -- set bridge sflow=@sflow
```

Parameter	Description	
SFLOW_AGENT	Indicates that the sFlow agent should send traffic from SFLOW_AGENT's IP address	
SFLOW_TARGET	Remote IP address of the sFLOW collector	
SFLOW_PORT	Remote IP destination port of the sFlow collector	
SFLOW_HEADE R	Size of packet header to sample (in bytes)	
SFLOW_SAMPLI NG	Sample rate	

To clear the sFLOW configuration, runt he following:

ovs-vsctl clear bridge br-vxlan mirrors

i) Note

Currently sFlow for OVS-DPDK is supported without CT.

CT CT NAT

To enable ct-ct-nat offloads in OvS-DPDK, execute the following command (deafult value is false):

```
ovs-vsctl set open_vswitch . other_config:ct-action-on-nat-
conns=true
```

If disabled, ct-ct-nat configurations will not be fully offloaded, improving connection offloading rate for

other cases (ct and ct-nat).

If enabled, ct-ct-nat configurations will be fully offloaded but ct and ct-nat offloading

will be slower to be created.

OpenFlow Meters (OpenFlow13+):

OpenFlow meters in OVS are implemented according to RFC 2697 (Single Rate Three Color Marker—srTCM).

- The srTCM meters an IP packet stream and marks its packets either green, yellow, or red. The color is decided on a Committed Information Rate (CIR) and two associated burst sizes, Committed Burst Size (CBS), and Excess Burst Size (EBS).
- A packet is marked green if it does not exceed the CBS, yellow if it exceeds the CBS but not the EBS, and red otherwise.
- The volume of green packets should never be smaller than the CIR.

To configure a meter in OVS:

- 1. Create a meter over a certain bridge:
 - 1. ovs-ofctl -0 openflow13 add-meter \$bridge
 meter=\$id,\$pktps/\$kbps,band=type=drop,rate=\$rate,
 [burst,burst_size=\$burst_size]

2. Parameters:

Par am ete r	Description	
bri dge	Name of the bridge on which the meter will be applied.	
id	Unique meter ID (32 bits) which will be used as an identifier for the meter.	

Par am ete r	Description	
pkt ps/ kbp s	Indication if the meter should work according to packets-per-second or kilobits-per-second.	
rat e	Rate of pktps/kbps of allowed data transmission.	
bur st	If set, enables burst support for meter bands through the "burst_size" parameter.	
bur st_ size	If burst is specified for the meter entry, configures the maximum burst allowed for the band in kilobits/packets, depending on whether kbps or pktps was specified. If unspecified, the switch is free to select some reasonable value depending on its configuration. Currently, if burst was not specified, the burst_size parameter is set as the "rate".	

2. Add the meter to a certain OpenFlow rule. For example:

ovs-ofctl -0 openflow13 add-flow \$bridge
"table=0,actions=meter:\$id,normal"

3. View the meter statistics:

ovs-ofctl -0 openflow13 meter-stats \$bridge meter=\$id

4. For more information, refer to openvswitch documentation <u>http://www.openvswitch.org/support/dist-docs/ovs-ofctl.8.txt</u>

VirtIO Acceleration through Hardware vDPA

Hardware vDPA Installation

Hardware vDPA requires QEMU v2.12 (or with upstream 6.1.0) and DPDK v20.11 as minimal versions.

To install QEMU:

1. Clone the sources:

```
git clone https://git.qemu.org/git/qemu.git
cd qemu
git checkout v2.12
```

2. Build QEMU:

```
mkdir bin
cd bin
../configure --target-list=x86_64-softmmu --enable-kvm
make -j24
```

To install DPDK:

1. Clone the sources:

git clone git://dpdk.org/dpdk cd dpdk git checkout v20.11

2. Install dependencies (if needed):

yum install cmake gcc libnl3-devel libudev-devel make
pkgconfig valgrind-devel pandoc libibverbs libmlx5 libmnl-

devel -y

3. Configure DPDK:

```
export RTE_SDK=$PWD
make config T=x86_64-native-linuxapp-gcc
cd build
sed -i 's/\(CONFIG_RTE_LIBRTE_MLX5_PMD=\)n/\1y/g' .config
sed -i 's/\(CONFIG_RTE_LIBRTE_MLX5_VDPA_PMD=\)n/\1y/g' .config
```

4. Build DPDK:

make -j

5. Build the vDPA application:

```
cd $RTE_SDK/examples/vdpa/
make -j
```

Hardware vDPA Configuration

To configure huge pages:

```
mkdir -p /hugepages
mount -t hugetlbfs hugetlbfs /hugepages
echo <more> > /sys/devices/system/node/node0/hugepages/hugepages-
1048576kB/nr_hugepages
```

```
echo <more> > /sys/devices/system/node/node1/hugepages/hugepages-
1048576kB/nr_hugepages
```

To configure a vDPA VirtIO interface in an existing VM's xml file (using libvirt):

1. Open the VM's configuration xml for editing:

virsh edit <domain name>

- 2. Modify/add the following:
 - 1. Change the top line to:

```
<domain type='kvm'
xmlns:qemu='http://libvirt.org/schemas/domain/qemu/1.0'>
```

2. Assign a memory amount and use 1GB page size for hugepages (size must be the same as used for the vDPA application), so that the memory configuration looks like the following.

```
<memory unit='KiB'>4194304</memory>
<currentMemory unit='KiB'>4194304</currentMemory>
<memoryBacking>
<hugepages>
<page size='1048576' unit='KiB'/>
</hugepages>
</memoryBacking>
```

3. Assign an amount of CPUs for the VM CPU configuration, so that the vcpu and cputune configuration looks like the following.

```
<vcpu placement='static'>5</vcpu>
<cputune>
<vcpupin vcpu='0' cpuset='14'/>
<vcpupin vcpu='1' cpuset='16'/>
<vcpupin vcpu='2' cpuset='18'/>
<vcpupin vcpu='3' cpuset='20'/>
<vcpupin vcpu='4' cpuset='22'/>
</cputune>
```

4. Set the memory access for the CPUs to be shared, so that the cpu configuration looks like the following.

```
<cpu mode='custom' match='exact' check='partial'>
   <model fallback='allow'>Skylake-Server-IBRS</model>
   <numa>
        <cell id='0' cpus='0-4' memory='8388608' unit='KiB'
memAccess='shared'/>
   </numa>
</cpu>
```

5. Set the emulator in use to be the one built in step "<u>2. Build QEMU</u>" above, so that the emulator configuration looks as follows.

<emulator><path to qemu executable></emulator>

6. Add a virtio interface using qemu command line argument entries, so that the new interface snippet looks as follows.

```
<qemu:commandline>
<qemu:arg value='-chardev'/>
<qemu:arg value='socket,id=charnet1,path=/tmp/sock-virtio0'/>
```

```
<qemu:arg value='-netdev'/>
<qemu:arg value='vhost-user,chardev=charnet1,queues=16,id=hostnet1'/>
<qemu:arg value='-device'/>
<qemu:arg value='virtio-net-
pci,mq=on,vectors=6,netdev=hostnet1,id=net1,mac=e4:11:c6:d3
page-per-vq=on,rx_queue_size=1024,tx_queue_size=1024'/>
</qemu:commandline>
```

Note: In this snippet, the vhostuser socket file path, the amount of queues, the MAC and the PCI slot of the VirtIO device can be configured.

Running Hardware vDPA

i Note

Hardware vDPA supports SwitchDev mode only.

Create the ASAP² environment:

- 1. Create the VFs.
- 2. Enter switchdev mode.
- 3. Set up OVS.

Run the vDPA application.

cd \$RTE_SDK/examples/vdpa/build
./vdpa -w <VF PCI BDF>, class=vdpa --log-level=pmd, info -- -i

Create a vDPA port via the vDPA application CLI.

```
create /tmp/sock-virtio0 <PCI DEVICE BDF>
```

Note: The vhostuser socket file path must be the one used when configuring the VM.

Start the VM.

```
virsh start <domain name>
```

For further information on the vDPA application, please visit: <u>https://doc.dpdk.org/guides/sample_app_ug/vdpa.html</u>.

Bridge Offload

i) Note

- Bridge offload is supported on ConnectX-6 Dx NIC
- Bridge offload is supported SwitchDev mode only
- Bridge offload is supported from kernel version 5.15

A Linux bridge is in-kernel software network switch (based on and implements subset of IEEE 802.1D standard) that is used to connect Ethernet segments together in a protocolindependent way. Packets are forwarded based on L2 Ethernet header addresses.

mlx5 provides capabilities to offload bridge data-plane unicast packet forwarding and VLAN management to hardware.

Basic Configuration

1. Initialize the ASAP² environment:

- 1. Create the VFs.
- 2. Enter switchdev mode.
- 2. Create a bridge and add mlx5 representors to bridge:

ip link add name bridge0 type bridge
ip link set enp8s0f0_0 master bridge0

Configuring VLAN

1. Enable VLAN filtering on the bridge.

ip link set bridge0 type bridge vlan_filtering 1

2. Configure port VLAN matching (trunk mode). In this configuration, only packets with specified VID are allowed.

bridge vlan add dev enp8s0f0_0 vid 2

3. Configure port VLAN tagging (access mode). In this configuration VLAN header in pushed/popped on reception/transmission on port.

bridge vlan add dev enp8s0f0_0 vid 2 pvid untagged

VF LAG Support

Bridge supports offloading on bond net device that is fully initialized with mlx5 uplink representors and is in single (shared) FDB LAG mode. Details about initialization of LAG are provided in <u>SR-IOV VF LAG</u> section, above.

Add bonding net device to bridge.

```
ip link set bond0 master bridge0
```

For further information on interacting with Linux bridge via iproute2 bridge tool, please consult man page (man 8 bridge).

Appendix: NVIDIA Firmware Tools

Download and install the MFT package corresponding to your computer's operating system. You would need the kernel-devel or kernel-headers RPM before the tools are built and installed.

The package is available at <u>nvidia.com/en-us/networking/</u> \rightarrow Products \rightarrow Software \rightarrow Firmware Tools.

1. Start the mst driver.

```
# mst start
Starting MST (Mellanox Software Tools) driver set
Loading MST PCI module - Success
Loading MST PCI configuration module - Success
Create devices
```

2. Show the devices status.

ST modules: -----MST PCI module loaded

```
MST PCI configuration module loaded
PCI devices:
_____
DEVICE_TYPE
                        MST
                                                       PCI
RDMA NET
                          NUMA
                          /dev/mst/mt4117_pciconf0.1
ConnectX4lx(rev:0)
                                                         04:00.1
net-enp4s0f1 NA
ConnectX4lx(rev:0) /dev/mst/mt4117_pciconf0
                                                         04:00.0
net-enp4s0f0 NA
# mlxconfig -d /dev/mst/mt4117_pciconf0 g | head -16
Device #1:
 _ _ _ _ _ _ _ _ _ _ _ _
Device type: ConnectX41x
PCI device: /dev/mst/mt4117_pciconf0
Configurations:
                                              Current
         SRIOV_EN
                                              True(1)
         NUM_OF_VFS
                                              8
         PF_LOG_BAR_SIZE
                                              5
         VF_LOG_BAR_SIZE
                                              5
         NUM_PF_MSIX
                                              63
         NUM_VF_MSIX
                                              11
         LINK_TYPE_P1
                                              ETH(2)
         LINK_TYPE_P2
                                              ETH(2)
```

- 3. Make sure your configuration is as follows:
 - * SR-IOV is enabled (SRIOV_EN=1)
 - * The number of enabled VFs is enough for your environment (NUM_OF_VFS=N)

* The port's link type is Ethernet (LINK_TYPE_P1/2=2) when applicable

If this is not the case, use mlxconfig to enable that, as follows:

1. Enable SR-IOV.

mlxconfig -d /dev/mst/mt4115_pciconf0 s SRIOV_EN=1

2. Set the number of required VFs.

mlxconfig -d /dev/mst/mt4115_pciconf0 s NUM_OF_VFS=8

3. Set the link type to Ethernet.

mlxconfig -d /dev/mst/mt4115_pciconf0 s LINK_TYPE_P1=2
mlxconfig -d /dev/mst/mt4115_pciconf0 s LINK_TYPE_P2=2

4. Conduct a cold reboot (or a firmware reset).

mlxfwreset -d /dev/mst/mt4115_pciconf0 reset

5. Query the firmware to make sure everything is set correctly.

mlxconfig -d /dev/mst/mt4115_pciconf0 q

Programming



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 <u>Raw Ethernet</u> <u>Programming: Basic Introduction—Code Example</u> 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 <u>Raw</u> <u>Ethernet Programming: Packet Pacing—Code Example</u> 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 <u>Raw Ethernet Programming:</u> <u>TSO—Code Example</u> 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 <u>Raw Ethernet Programming: ToS—Code Example</u> 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 <u>Raw Ethernet Programming: Flow ID Steering</u>—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 <u>Raw Ethernet Programming: VXLAN Steering—Code Example</u> Community post.

Device Memory Programming

(i) 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 on 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

j) 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 '-I'):

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

Utili ty	Description	
du mp _fts	Dumps tables for every switch found in an ibnetdiscover scan of the subnet. The dump file format is compatible with loading into OpenSM using the -R file -U /path/to/dump-file syntax. For further information, please refer to the tool's man page.	
iba ddr	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.	
ibc ach eed it	Allows users to edit an ibnetdiscover cache created through thecache option in ibnetdiscover(8). For further information, please refer to the tool's man page.	
ibc cco nfig	Supports the configuration of congestion control settings on switches and HCAs. For further information, please refer to the tool's man page.	
ibc cqu ery	Supports the querying of settings and other information related to congestion control. For further information, please refer to the tool's man page.	

Utili ty	Description			
ibc ong est	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 "/opt/ ibutils2/bin/ibcongest -h". For further information, please refer to the tool's man page.			
ibd ev2 net dev	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.			
ibdi agn et (of ibut ils2)	 For further information, please refer to the tool's man page. 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: Fabric discovery Duplicated GUIDs detection Links in INIT state and unresponsive links detection Counters fetch Error counters check Routing checks Link width and speed checks Alias GUIDs check Subpet Mapager check 			
ibdi agp ath	nodes and ports traversed along the path. It utilizes device specific health queries			

Utili ty	Description			
	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. Example: ibdiagpathsrc_lid 1dest_lid 28 For further information, please refer to the tool's -help flag.			
	Dump InfiniBand traffic that flows to and from NVIDIA's ConnectX® family adapters InfiniBand ports. Note the following:			
	ibdump is not supported for Virtual functions (SR-IOV)Infiniband traffic sniffing is supported on all HCAs			
	The dump file can be loaded by the Wireshark tool for graphical traffic analysis. The following describes a workflow for local HCA (adapter) sniffing:			
ibd um p	 Run ibdump with the desired options Run the application that you wish its traffic to be analyzed Stop ibdump (CTRL-C) or wait for the data buffer to fill (inmem-mode) Open Wireshark and load the generated file 			
	To download Wireshark for a Linux or Windows environment go to <u>www.wireshark.org</u> . Notes:			
	 Although ibdump is a Linux application, the generated .pcap file may be analyzed on either operating system. If one of the HCA's ports is configured as InfiniBand, ibdump requires IPoIB DMFS to be enabled. For further information, please refer to <u>Flow Steering Configuration</u> section. 			
	For further information, please refer to the tool's man page.			
ibli nki nfo	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. For further information, please refer to the tool's man page.			
ibn etdi sco ver	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). 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. For further information, please refer to the tool's man page.			

Utili ty	Description		
ibn ets plit	Automatically groups hosts and creates scripts that can be run in order to split the network into sub-networks containing one group of hosts. For further information, please refer to the tool's man page.		
ibn ode s	Uses the current InfiniBand subnet topology or an already saved topology file and extracts the InfiniBand nodes (CAs and switches). For further information, please refer to the tool's man page.		
ibpi ng	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. For further information, please refer to the tool's man page.		
	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:		
ibp ort	disable, enable or reset the portvalidate the port's link width and speed against the peer port		
sta te	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:		
	The first ACTIVE port that is found.If not found, the first port that is UP (physical link state is LinkUp).		
	For further information, please refer to the tool's man page.		
ibq uer	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.		
yerr ors	In addition to reporting errors on every port, ibqueryerrors can report the port transmit and receive data as well as report full link information to the remote port if available. For further information, please refer to the tool's man page.		
ibro ute	Uses SMPs to display the forwarding tables—unicast (LinearForwarding- Table or LFT) or multicast (MulticastForwardingTable or MFT)—for the specified switch LID and the optional lid (mlid) range. The default range is all valid entries in the range 1 to FDBTop.		
ibst at	For further information, please refer to the tool's man page. ibstat 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. For further information, please refer to the tool's man page.		

Utili ty	Description		
ibst atu s	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.		
ibs wit che s	Traces the InfiniBand subnet topology or uses an already saved topology file to extract the InfiniBand switches. For further information, please refer to the tool's man page.		
ibsy sst at	Uses vendor mads to validate connectivity between InfiniBand nodes and obtain other information about the InfiniBand node. ibsysstat is run as client/ server. The default is to run as client. For further information, please refer to the tool's man page.		
ibto pod iff	 Compares a topology file and a discovered listing of subnet.lst/ibdiagnet.lst and reports mismatches. Two different algorithms provided: Using the -e option is more suitable for MANY mismatches it applies less heuristics and provide details about the match Providing the -s, -p and -g starts a detailed heuristics that should be used when only small number of changes are expected For further information, please refer to the tool's man page. 		
ibtr ace rt	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. For further information, please refer to the tool's man page.		
ibv_ asy nc wat ch	Display asynchronous events forwarded to userspace for an InfiniBand device. For further information, please refer to the tool's man page.		
ibv_ dev ices	Lists InfiniBand devices available for use from userspace, including node GUIDs. For further information, please refer to the tool's man page.		
ibv_ dev info	Queries InfiniBand devices and prints about them information that is available for use from userspace. For further information, please refer to the tool's man page.		

Utili ty	Description		
mst flin t	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. To run mstflint, you must know the device location on the PCI bus. Note: If you purchased a standard NVIDIA network adapter card, please download the firmware image from <u>nvidia.com/en-us/networking/</u> \rightarrow Support \rightarrow Support \rightarrow Firmware Download. If you purchased a non-standard card from a vendor other than NVIDIA, please contact your vendor. For further information, please refer to the tool's man page.		
per fqu ery	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. For further information, please refer to the tool's man page.		
saq uer y	Issues the selected SA query. Node records are queried by default. For further information, please refer to the tool's man page.		
smi nfo	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. Note: 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.		
sm par que ry	Sends SMP query for adaptive routing and private LFT features. For further information, please refer to the tool's man page.		
sm pdu mp	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.		
sm pqu ery	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/<u>ibdiagnet2.pm</u>.

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.

Utili ty	Description		
ib_a tom ic_b w	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.		
ib_a tom ic_l at	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. For further information, please refer to the tool's man page.		
ib_r ead _bw	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. Read is available only in RC connection mode (as specified in IB spec). For further information, please refer to the tool's man page.		
ib_r ead _lat	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). For further information, please refer to the tool's man page.		
ib_s end _bw	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.		
ib_s end _lat	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		

Utili ty	Description		
	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.		
ib_ writ e_b w	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 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.		
ib_ writ e_la t	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.		
raw _et her net _bw	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.		
raw _et her net _lat	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 <u>networking-support@nvidia.com</u>.

The chapter contains the following sections:

- General Issues
- Ethernet Related Issues
- InfiniBand Related Issues
- Installation Related Issues
- <u>Performance Related Issues</u>
- SR-IOV Related Issues
- <u>PXE (FlexBoot) Related Issues</u>
- RDMA Related Issues
- <u>Debugging Related Issues</u>
- OVS Offload Using ASAP2 Direct Related Issues

General Issues

Issue	Cause	Solution
The system panics when it is booted with a failed adapter installed.	Malfunction hardware component	 Remove the failed adapter. Reboot the system.
NVIDIA adapter is not identified as a PCI device.	PCI slot or adapter PCI connector dysfunctionality	 Run Ispci. Reseat the adapter in its PCI slot or insert the adapter to a different PCI slot. If the PCI slot confirmed to be functional, the adapter should be replaced.
NVIDIA adapters are not installed in the	Misidentification of the NVIDIA adapter installed	Run the command below and check NVIDIA's MAC to

Issue	Cause	Solution
system.		identify the NVIDIA adapter installed. Ispci grep Mellanox' or 'Ispci -d 15b3: Note : NVIDIA MACs start with: 00:02:C9:xx:xx:xx, 00:25:8B:xx:xx:xx or F4:52:14:xx:xx:"
The default device may vary when invoking user apps (such as ibv_asyncwatch) which run using a specific device.	The default device for such apps is the first device in the device list generated by libibverbs. This first device in the list varies, depending on which and how many InfiniBand devices are installed on the host, which slot the devices are installed on, whether they use SR-IOV, and other factors.	Always specify the desired device explicitly when running userspace apps, by using the provided command line parameter (for example: ibv_asyncwatch -d <dev>).</dev>
Insufficient memory to be used by udev upon OS boot.	udev is designed to fork() new process for each event it receives so it could handle many events in parallel, and each udev instance consumes some RAM memory.	Limit the udev instances running simultaneously per boot by adding udev.children- max= <number> to the kernel command line in grub.</number>
Operating system running from root file system located on a remote storage (over NVIDIA devices), hang during reboot/shutdown (errors such as "No such file or directory" will appear).	The openibd service script is called using the 'stop' option by the operating system. This option unloads the driver stack. Therefore, the OS root file system disappears before the reboot/ shutdown procedure is completed, leaving the OS in a hang state.	Disable the openibd 'stop' option by setting 'ALLOW_STOP=no' in /etc/ infiniband/openib.conf configuration file.
NVIDIA adapter warning print to dmesg: Detected insufficient power on the PCIe slot (xxxW).	Insufficient PCI power.	Investigate the cause for lack of PCI power.

Ethernet Related Issues

Issue	Cause	Solution	
Ethernet interfaces renaming fails leaving them with names such as renameXY.	Invalid udev rules.	Review the udev rules inside the "/etc/ udev/rules.d/70-persistent-net.rules" file. Modify the rules such that every rule is unique to the target interface, by adding correct unique attribute values to each interface, such as dev_id, dev_port and KERNELS or address). Example of valid udev rules: SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{dev_id}=="0x0", ATTR{type}=="1", KERNEL=="eth*", ATTR{dev_port}=="0", KER- NELS=="0000:08:00.0", NAME="eth4" SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{dev_port}=="0x0", ATTR{type}=="1", KERNEL=="eth*", ATTR{dev_id}=="0x0", ATTR{type}=="1", KERNEL=="eth*", ATTR{dev_id}=="0x0", ATTR{type}=="1", KERNEL=="eth*",	
No link.	Misconfigurati on of the switch port or using a cable not supporting link rate.	 Ensure the switch port is not down Ensure the switch port rate is configured to the same rate as the adapter's port 	
Degraded performance is measured when having a mixed rate environment (10GbE, 40GbE and 56GbE).	Sending traffic from a node with a higher rate to a node with lower rate.	 Enable Flow Control on both switch ports and nodes: On the server side run: ethtool -A <interface> rx on tx on</interface> On the switch side run the following command on the relevant interface: send on force and receive on force 	
No link with break- out cable.	Misuse of the break-out cable or misconfigurati on of the switch's split ports	 Use supported ports on the switch with proper configuration. For further information, please refer to the MLNX_OS User Manual. Make sure the QSFP breakout cable side is connected to the SwitchX. 	

InfiniBand Related Issues

Issue	Cause	Solution
The following messages is logged after loading the driver: multicast join failed with status - 22	Trying to join a multicast group that does not exist or exceeding the number of multicast groups supported by the SM.	If this message is logged often, check for the multicast group's join requirements as the node might not meet them. Note: If this message is logged after driver load, it may safely be ignored.
Unable to stop the driver with the following on screen message: ERROR: Module <module> is in use</module>	An external application is using the reported module.	Manually unloading the module using the 'modprobe -r' command.
Logical link fails to come up while port logi- cal state is Initializing.	The logical port state is in the Initializing state while pending the SM for the LID assignment.	 Verify an SM is running in the fabric. Run 'sminfo' from any host connected to the fabric. If SM is not running, activate the SM on a node or on managed switch.
InfiniBand utilities commands fail to find devices on the system. For example, the 'ibv_devinfo' command fail with the following output: Failed to get IB devices list: Function not implemented	The InfiniBand utilities commands are invoked when the driver is not loaded.	Load the driver: /etc/init.d/openibd start

Installation Related Issues

Installation Issues

Issue	Cause	Solution
Driver installation fails.	The install script may fail for the following reasons:	 Use only supported installation options. The full list of installation

Issue	Cause	Solution
	 Using an unsupported installation option Failed to uninstall the previous installation due to dependencies being used The operating system is not supported The kernel is not supported. You can run mlnx_add_kernel_sup- port.sh in order to to generate a MLNX-OFED package with drivers for the kernel Required packages for installing the driver are missing Missing kernel backport supported kernel 	 options case be displayed on screen by using: mlnxofedinstallh Run 'rpm -e' to display a list of all RPMs and then manually uninstall them if the preliminary uninstallation failed due to dependencies being used. Use a supported operating system and kernel Manually install the missing packages listed on screen by the installation failed due to missing prerequisites.
After driver installation, the openibd service fail to start. This message is logged by the driver: Unknown symbol	The driver was installed on top of an existing In-box driver.	 Uninstall the MLNX_OFED driver. ofed_uninstall.sh Reboot the server. Search for any remaining installed driver. If found, move them to the /tmp directory from the current directory. Re-install the MLNX_OFED driver. Restart the openibd service.

Fixing Application Binary Interface (ABI) Incompatibility with MLNX_OFED Kernel Modules

(i) Note

This section is relevant for RedHat and SLES distributions only.

Overview

MLNX_OFED package for RedHat comes with RPMs that support KMP (weak-modules), meaning that when a new errata kernel is installed, compatibility links will be created under the weak-updates directory for the new kernel. Those links allow using the existing MLNX_OFED kernel modules without the need for recompilation. However, at times, the ABI of the new kernel may not be compatible with the MLNX_OFED modules, which will prevent loading them. In this case, the MLNX_OFED modules must be rebuilt against the new kernel.

Detecting ABI Incompatibility with MLNX_OFED Modules

When MLNX_OFED modules are not compatible with a new kernel from a new OS or errata kernel, no links will be created under the weak-updates directory for the new kernel, causing the driver load to fail. Checking for the existence of needed module links under weak-updates directory can be done by reloading the MLNX_OFED modules. If one or more modules are missing, the driver reload will fail with an error message.

Example:

Resolving ABI Incompatibility with MLNX_OFED Modules

In order to fix ABI incompatibility with MLNX_OFED modules, the modules should be recompiled against the new kernel, using the mlnx_add_kernel_support.sh script, available in MLNX_OFED installation image.

There are two ways to recompile the MLNX_OFED modules:

1. Local recompilation and installation on one server.

Run the mInxofedinstall command to recompile the kernel modules and reinstall the whole MLNX_OFED on the server. Mount MLNX_OFED ISO image or extract the TGZ file:

cd <MLNX_OFED dir>
./mlnxofedinstall --skip-distro-check --add-kernel-support
--kmp --force

Notes:

- The --kmp flag will enable rebuilding RPMs with KMP (weak-updates) support for the new kernel. Therefore, in the next OS/kernel update, the same modules can be used with the new kernel (assuming that the ABI compatibility was not broken again).

- The command above will rebuild only the kernel RPMs (using mlnx_add_kernel_support.sh), and will save the resulting MLNX_OFED package under /tmp and start installing it automatically. This package can be used for installation on other servers using regular mlnxofedinstall command or yum.

- 2. Preparing a new image on one server and deploying it on the cluster.
 - 1. Use the mlnx_add_kernel_support.sh script directly only to rebuild the kernel RPMs (without running any installations) on one server. Mount MLNX_OFED ISO image or extract the TGZ file:

cd <MLNX_OFED dir>

Note: This command will save the resulting MLNX_OFED package under /tmp.

Example:

2. Install the newly created MLNX_OFED package on the cluster:

Option 1: Copy the package to the servers and install it using the mInxofedinstall script.

Option 2: Deploy the MLNX_OFED package using YUM (for YUM installation instructions, refer to <u>Installing MLNX_OFED Using YUM</u> section):

i. Extract the resulting MLNX_OFED image and copy it to a shared NFS location.

ii. Create a YUM repository configuration.

iii. Install the new MLNX_OFED kernel RPMs on the servers: # yum update **Example**:

```
. . .
. . .
_______
Package
                   Arch Version
Repository Size
______
Updating:
epel-release
                  noarch 7-7
epel
    14 k
kmod-iser
                   x86_64 1.8.0-
OFED.3.3.1.0.0.1.gf583963.201606210906.rhel7u1
mlnx_ofed 35 k
kmod-isert
                   x86_64 1.0-
OFED.3.3.1.0.0.1.gf583963.201606210906.rhel7u1
mlnx ofed 32 k
kmod-kernel-mft-mlnx x86_64 4.4.0-
1.201606210906.rhel7u1
mlnx ofed 10 k
kmod-knem-mlnx
                  x86_64 1.1.2.90mlnx1-
OFED.3.3.0.0.1.0.3.1.ga04469b.201606210906.rhel7u1 mlnx_ofed
22 k
kmod-mlnx-ofa_kernel
                x86_64
                         3.3-
OFED.3.3.1.0.0.1.gf583963.201606210906.rhel7u1
mlnx_ofed 1.4 M
                   x86_64 1.6.0-
kmod-srp
OFED.3.3.1.0.0.1.gf583963.201606210906.rhel7u1
mlnx_ofed 39 k
Transaction Summary
______
Upgrade 7 Packages
```



Note: The MLNX_OFED user-space packages will not change; only the kernel RPMs will be updated. However, "YUM update" can also update other inbox packages (not related to OFED). In order to install the MLNX_OFED kernel RPMs only, make sure to run:

```
# yum install mlnx-ofed-kernel-only
```

Note: mlnx-ofed-kernel-only is a metadata RPM that requires the MLNX_OFED kernel RPMs only.

3. Verify that the driver can be reloaded:

/etc/init.d/openibd restart

Performance Related Issues

Issue	Cause	Solution
The driver works but the transmit and/or receive data rates are not optimal.	-	These recommendations may assist with gaining immediate improvement: 1. Confirm PCI link negotiated uses its maximum capability 2. Stop the IRQ Balancer service: /etc/init.d/irq_balancer stop 3. Start mInx_affinity service: mInx_affinity start

Issue	Cause	Solution
		For best performance practices, please refer to the "Performance Tuning Guide for NVIDIA Network Adapters".
Out of the box throughput performance in Ubuntu14.04 is not optimal and may achieve results below the line rate in 40GE link speed.	IRQ affinity is not set properly by the irq_balancer	For additional performance tuning, please refer to Performance Tuning Guide.

SR-IOV Related Issues

Issue	Cause	Solution
When assigning a VF to a VM the following message is reported on the screen: PCI-assgine: error: requires KVM support	SR-IOV and virtualization are not enabled in the BIOS.	 Verify they are both enabled in the BIOS Add to the GRUB configuration file to the following kernel parameter: "intel_immun=on" (see "<u>Setting Up SR-IOV</u>" section).

PXE (FlexBoot) Related Issues

Issue	Cause	Solution
PXE boot timeout.	The 'always-broadcast' option is disabled.	Enable 'always-broadcast on'. For the complete procedure, please refer to <u>Linux PXE User</u> <u>Guide</u> .
PXE InfiniBand link fails with the following messages although the DHCP request was sent:	Either the SM is not running in the fabric or the SM default multicast group was created with non-default settings.	 Activate the SM on a node or on managed switch. Check in the SM partitions.conf file that the default partition rate and

Issue	Cause	Solution
<i>Initializing</i> and <i>The</i> <i>socket is not</i> <i>connected</i> .		MTU setting are SDR and 2K, respectively. The PXE is establishing by default an SDR link set with an MTU of 2K. If the default multicast group opened with different rate and/or MTU, the SM will deny the PXE request to join.
NVIDIA adapter is not identified as a boot device.	The expansion ROM image is not installed on the adapter. or the server's BIOS is not configured to work on Legacy mode	 Run a flint query to display the expansion ROM information. For example: "flint -d /dev/mst/ mt4099_pci_cr0 q" and look for the "Rom info:" line. For further information on how to burn the ROM, please refer to MFT User Manual. Make sure the BIOS is configured to work in Legacy mode if the adapter's firmware does not include a UEFI image.

RDMA Related Issues

Issue	Cause	Solution
Infiniband-diags tests, such as	When running a test between 2	Run the test using
'ib_write_bw', fail between	systems in the fabric with	the same perftest
systems with different driver	different Infiniband-diags	RPM on both
releases.	packages installed.	systems.

Debugging Related Issues

Issue	Cause	Solution
False positive errors when running applications with valgrind.	Default MLNX_OFED libraries are compiled with- out valgrind support and several resources are managed by the kernel.	Libraries' files compiled with valgrind support are installed under "/usr/ lib64/mlnx_ofed/valgrind/" • To run an application over these libraries, thus prevent false positive errors: # env LD_LIBRARY_PATH=/usr/ lib64/mlnx_ofed/valgrind/ val- grind [valgrind options] <application cmd=""> • To suppress most of valgrind's false positive errors, generate the suppression file: #./generate_mlnx_ofed supp.sh > mlnx.supp</application>

OVS Offload Using ASAP2 Direct Related Issues

Issue	Cause	Solution(s)
Traffic is not offloaded	 OVS uses TC flower classifier to add offloading rules to both the software and the hardware. TC flower classifier fails to add a rule. A rule was added to the TC flower classifier but failed to be added to the firmware. 	 Check for system error in dmesg or the system logging facility like journalctl Check OVS logs for errors Dump the rules using the TC command line For example: Dump rules on a specific interface # tc filter show dev ens4f0 parent ffff:

Common Abbreviations and Related Documents

Common Abbreviations and Acronyms

Abbreviation/ Acronym	Description
В	(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
НСА	Host Channel Adapter
HW	Hardware
IB	InfiniBand
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*.

Term	Description
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

Term	Description
Subnet Manager (SM)	One of several entities involved in the configuration and control of the IB fabric
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 <u>Release Notes</u> 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

- <u>Release Notes History</u>
- User Manual Revision History

Release Notes History

- Changes and New Features History
- <u>Bug Fixes History</u>

Changes and New Features History

This section includes history of changes and new feature of three major (GA) releases back. For older versions' history, please refer to their dedicated release notes.

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

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

Feature/Change	Description		
	24.04-0.7.0.0		
General	Bug fixes		
	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 <u>SR-IOV Live Migration</u> .		
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: Added XDP multi-buffer support to the default RQ type (Striding RQ) Added support for XDP multi-buffer redirect-in Allowed non-linear single-segment frames in XDP TX MPWQE 		
DC MRA Caps	Increased the maximum outstanding DC atomic reads to 32 to enable the driver to handle reads larger than 8K. Now log_max_ra_{res/req}_dc is exposed in DV API, enabling the user to see both RC and DC values.		

Feature/Change	Description	
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 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. \$ 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	
Dynamic EQ	When an EQ is not created, the default affinity value is empty, therefore the "show" command output will show an empty affinity.	
General	General <u>Bug fixes</u>	
24.01-0.3.3.1		
Installation		
Kernel Version Update for Kylin 10	Upgraded the kernel version of the Kylin 10 Operating System Service Pack 3 to v4.19.90-52.22.v2207.	
General		
General	Bug fixes	

Feature/Change	Description	
23.10-0.5.5.0		
ASAP ² Features		
Rate Limiter Extension for LAG Mode	[ConnectX-5 and above] During rate configuration, this feature considers the aggregated port speed when LAG is configured, so the functions (PF,VF,SF) can utilize the maximum aggregated link speed for transmission.	
Core Features		
Live Migration Support with	[ConnectX-7 and above] Added an option to perform live migration on a machine with IPsec enabled. This allows having a VF on VM that	

Feature/Change	Description		
IPsec Full Offload	has a valid IPsec state and migrate it without having to remove the IPsec. Note: This feature is currently not supported when working in MPV mode (Multi port VHCA/Dual port). Such support will be added in future releases.		
Exposing the Max SF Configuration via Devlink Resource	[All HCAs] This feature allows user applications to view the maximum number of SFs (Scalable Functions) configured on the system via general Linux API, Devlink resource.		
NetDev Features			
NIC Temperature Exposure	[ConnectX-5 and above] Added support for NIC temperature exposure by implementing the hwmon kernel API.		
Counters Indicating Packet Drops due to Catastrophic Steering Errors	[All HCAs] Added two counters through the Devlink health reporter - generated_pkt_steering_fail and handled_pkt_steering_fail. These counters indicate packet drops due to catastrophic steering errors.		
RDMA Features			
Using the RDMA Tool to Allow Controlled QKEYs Use by Non- privileged Users	[All HCAs] To allow using controlled QKEYs by non-privileged users, the following command can now be used in the iproute2 RDMA tool, instead of the "enforce_qkey_check" module parameter, which will be deprecated: \$rdma system set privileged-qkey on		
IPsec Configuration Support When All IPsec Operations are Offloaded to the Hardware	[ConnectX-7 and above] Added support for configuring IPsec even when all IPsec operations are completely offloaded to the hardware. This feature not only provides a significant performance improvement, but also enables the use of IPsec over RoCE packets, which are outside the network stack and cannot be used without full hardware offload. As a result, users can now leverage the benefits of IPsec protocol with RoCE V2, even when using SR-IOV VFs.		
Security Features	Security Features		
Configuration from VM/VF for Full Offload Transport Mode	[ConnectX-6 Lx and above] This feature introduces two new boolean attributes of a port function: ipsec_crypto and ipsec_packet (also known as IPsec Full Offload). The goal is to provide a level of granularity for controlling VF IPsec offload capabilities similar to what is currently offered in the software mode. This allows users to decide if they want both types of IPsec offload enabled for a VF, just one of them, or none (which is the default).		

Feature/Change	Description
General	
General	Bug fixes

Feature/Change	Description		
23.07-0.5.1.2	23.07-0.5.1.2		
RDMA Features			
QKEY Mitigation in the Kernel	 QKEY creation with the MSB set is available now for non-privileged users as well. To allow non-privileged users to create QKEY with MSB set, the below new module parameter was added to ib_uverbs module: Module Parameter: enforce_qkey_check Description: Force QKEY MSB check for non-privileged user on UD QP creation Default: 0 (disabled) Note: In this release, this module parameter is disabled by default to ensure backward compatibility and give customers the opportunity		
	to update their applications accordingly. In the upcoming release, it will be enabled by default, and later on deprecated.		
	23.07-0.5.0.0		
ASAP ² Features			
Header Rewrite and CT Support	[BlueField-2] Added support for offloading CT rules with header rewrite of L3.		
CT UDP Unidirectional Traffic Offload	[All HCAs] Added support for offloading long-running unidirectional UDP connection with Conntrack.		
CT Rules with NAT and Mirror Offload	[BlueField-2] Added support for offloading CT rules with NAT and mirror.		
Core Features			
Embedded CPU Virtual Functions	[Bluefield-2] Added the ability to create virtual functions within the ARM.		
Full Chip Reset on BlueField-2	[BlueField-2] Full chip reset in DPU mode is now supported on BlueField-2 when using mlxfwreset with the "sync 1" option. During this flow, the ARM is rebooted and the firmware is reloaded.		

Feature/Change	Description		
Firmware Page Limit via VHCA_ICM_CTRL	[ConnectX-5 and above] Added the ability to expose page limit via the VHCA_ICM_CTRL FW register. The expected behavior is similar to the previous page limit, except for the firmware size limitation. The new max limit is 2^32-1, which represents 'no limit'.		
PRE_COPY State Support in the mlx5-vfio Driver [ConnectX-6 Dx and above] Added support for PRE_COPY m The optional PRE_COPY state opens the saving data transfe before reaching STOP_COPY, and thereby helps reducing the downtime of the VM.			
Lightweight Local SFs	[ConnectX-5 and above] Added a feature that decreases the amount of time needed to probe and configure SFs (s ub-functions) by reducing the time of Devlink reload of the SF.		
QEMU Live Migration Support	[ConnectX-6 Dx and above] Added support for QEMU VM migration with an assigned VF from one source host to another destination host. This is achieved as part of the general QEMU migration flow, which involves suspending the VF on the source host, transferring all its data to the destination host, and resuming the VF on the destination. Note: This migration feature includes basic functionality only, and does not yet support advanced features such as dirty page tracking or pre-copy.		
4 Ports VF LAG Support	[ConnectX-7] Added support for VF LAG over 4 ports HCAs. Note: Only LAGs with all HCA ports are supported. Meaning, with a 4-port HCA, only 4 port LAG is supported. A 2-port or a 3-port LAG is not supported.		
Installation	Installation		
FIPS Compliance	 [All HCAs] Starting from v23.07, OFED is FIPS (Federal Information Processing Standard) compliant for the RH8.8, RH9.2 and SLES15.5 operating systems. The following limitations should be considered: OFED Installation with FIPS support requires a manual firmware upgrade. OFED Installation with add-kernel-support will result in the removal of the FIPS support. 		
NetDev Features			
Enhanced CQE Compression	[ConnectX-5 and above] Added support for the enhanced version of the RX CQE compression hardware feature. By compressing the RX CQEs, the PCI bandwidth utilization is improved and the load on it is		

Feature/Change	Description
	reduced. The enhanced version of this device feature has improved latency and CPU utilization.
256 Bit Keys with kTLS Offload Support	[BlueField-2] Added support for kTLS offload with a 256 bit key size.
Improving Affinity Hints according to Numa Distances	[All HCAs] Updated the binary NUMA preference for the system to consider actual distances. Following the update, remote NUMA nodes with shorter distances are preferred over further ones, rather than relying solely on local/remote distinctions.
Flow Steering Decoupling	[All HCAs] Added the ability to decouple flow steering into a separate module, making it loosely coupled and thus easier to read, maintain and debug.
RDMA Features	
IPv6 Address Support	[All HCAs] Reduced the storage IO latency that occurred when establishing a large number of RDMACM (RDMA Connection Manager) connections by setting the RDMACM RoCE static rate to 0.
Reducing Storage IO Latency for a Large Number of RDMACM Connections	[All HCAs] Reduced the storage IO latency that occurred while establishing a large number of RDMACM (RDMA Connection Manager) connections by setting the RDMACM RoCE static rate to 0.
QKEY Mitigation in the Kernel	 [All HCAs] Non-privileged users are now blocked from setting controlled/privileged QKEYs (QKEY with MSB set). To allow non-privileged users to create a QKEY with an MSB set, a new module parameter was added to the ib_uverbs module: module parameter: enforce_qkey_check Description: Force a QKEY MSB check for non-privileged users on UD QP creation Default is 0 (disabled) In this release, this module parameter is introduced in disabled state in order to ensure backward compatibility and allow the required applications updates. In the upcoming release, this feature will be enabled by default, and later on deprecated.
Security	1
Fast Update Encryption Key	[ConnectX-6 Dx and above] Enhanced the performance of DEK operations by introducing a DEK pool that serves the same key

Feature/Change	Description	
Support	purpose and utilizes bulk allocation, destruction and invalidation provided by the firmware. Users can now retrieve a DEK object from the pool, and update it with a key using the modify_DEK command.	
Software Steering		
Matching over BTH Acknowledge Bit Support	[Connect-6 DX, Connect-7 and BlueField-2] Added support for matching over BTH acknowledge bit using the mlx5dv_dr API.	
General	Bug fixes	

Feature/Ch ange	Description
	23.04-1.1.3.0
Operating Systems	Added support for the following operating systems: • RHEL 9.2 • RHEL 8.8 • SLES15-SP5
23.04-0.5.3.3	3
ASAP ² Featu	Jres
VXLAN GBP Options Offload with OVS	[ConnectX-6 Dx and above] Added support of encap and decap of VXLAN tunnel with GBP options offload with OVS.
TC Rules: Additional Actions	[ConnectX-5 and above] Added support for adding TC rules with trap action with additional actions (mirror and pedit).
Multiport E-Switch	[ConnectX-6 Dx] Added support for Multiport E-Switch, a mode where a single E-Switch connects all VPorts and physical ports on the NIC. This allows for scenarios such as sending traffic from a VF created on PFO to an uplink that is natively associated with the uplink of PF1.
OVS Offload with MACVLAN Interface	[ConnectX-6 Dx and above] Added support for OVS offload when MACVLAN interface above bond is attached to OVS bridge.

Feature/Ch ange	Description
Above Bond	
Core Feature	es
PCC fwtrace	[ConnectX-6 Dx and BlueField-2] Added support for installing a special user image into the firmware, which can be burned into either of two available slots for such applications, which enables monitoring the image's activities using the fw_tracer located inside the mlx5 driver. To view the output of the tracer, the user can access trace point, but it is important to note that they can only view traces that are generated after enabling the trace point by using the following command: 'echo 1 > sys/kernel/debug/tracing/events/mlx5/mlx5_fw/enable'." Note: Not supported in ConnectX-6 Lx adapter cards.
Relaxed Ordering in VFs/QEMU Assigned VFs	[ConnectX-6 and above] Added support for using Relaxed Ordering in VFs directly and in VFs assigned to QEMU. Relaxed Ordering can significantly improve performance on certain setups and, until now, it could be used only in PFs.
Migratable Bit	[ConnectX-6 Dx and above] Added support for Migratable Bit in Live Migration. Because some features cannot be migrated, such as IPsec, for example, when VF is marked as migratable, those features are disabled. This feature allows the user to configure whether VF can be migrated.
Live Migration Dirty Page Tracking Support in Linux	[ConnectX-7] The dirty pages tracking support enables reducing downtime upon live migration. Once it is used, only the pages that were really dirtied by the device will be marked in QEMU as dirty and will be sent to the target upon stop. Without dirty tracking, all RAM is marked dirty so all RAM is resent upon stop and the downtime is increased.
NetDev Feat	ures
Configurin g Hairpin Queue Size	[ConnectX-5 and above] Added the ability to configure the number and size of hairpin queues through devlink param command. For example: devlink dev param set pci/0000:08:00.0 name hairpin_queue_size value 512 cmode driverinit devlink dev param set pci/0000:08:00.0 name hairpin_num_queues value 1 cmode driverinit
RDMA Features	

Feature/Ch ange	Description		
CC - RTT response SL (CNP SL)	[ConnectX-6 Dx and above] Users can now customize the DSCP value of RTT Response packets (in Ethernet, using debugfs). This feature allows for prioritization of RTT Response packets, preventing any delay that might lead to incorrect congestion assumptions on the RTT Requester side.		
Expose VF RoCE Statistics on Host Side (Represent or)	[All HCAs] Added support for allowing the host to track various statistics for the VFs, specifically all of the Q_counters stats even if he gives the VF to a VM, through the representor Q_counters for the VF which are now exposed over the host.		
Fatal QP Error Logging	[All HCAs] With this feature, a kernel error log is now generated when certain fatal QP errors occur.		
Enabling Selective Repeat by Default	[ConnectX-7 and above] Added support for making Selective Repeat (SR) protocol to be anabled by default. SR retransmits only the frame that is damaged and not all of the frames that were sent. As such, SR makes more efficient use of network bandwidth.		
Software Ste	Software Steering		
Matching IB BTH in RoCE Packets	[ConnectX-6 Dx and above] Added support for matching the IB BTH in RoCE packets. One example of this match is that the user can monitor RoCEv2 CNP (Congestion Notification Packet) by matching BTH opcode 0x81.		
General	Bug fixes		

Feature/Ch ange	Description		
5.9-0.5.6.0	5.9-0.5.6.0		
ASAP ² Featu	ASAP ² Features		
Linux Bridge VLAN Filtering of 802.1 Q Packets	[ConnectX-6 Dx] Extended mlx5 Linux bridge VLAN offload to support packets tagged with 802.1 Q VLAN ethertype.		

Feature/Ch ange	Description	
Offloading sFlow Sampling Rules	[ConnectX-5 and above] Added support for sFlow sampling rules offloads. sFlow is an industry standard technology for monitoring high speed switched networks. Open vSwitch integrated sFlow to extend the visibility into virtual servers, ensuring data center visibility and control.	
Core Features		
Configuring Shared Buffer Size	As with each port buffer command the user triggers, the shared buffer	
Control SF Class	F [All HCAs] Added support for Control SF Class. Each PCI, PF, VF, SF function, by default, has netdevice, RDMA, and vdpa-net devices always enabled. This feature enables the user to control which device functionality to enable/disable. Note: Requires kernel 5.18 or higher.	
Installation Features		
	[All HCAs] Added support for ip2gid tool. This tool does the following:1. Resolves a destination IP into a destination GID needed when running a rdmacm applications (ip2gid).	
ip2gid Tool	2. Resolves a GID into one PR (PathRecord) or multiple PRs if needed (gid2lid).	
	This tool is needed when rdmacm is used to initiate InfiniBand traffic between nodes on different IP subnets in InfiniBand fabrics.	
NetDev Featu	ures	
Support RSS over XSK Queues	[All HCAs] Use default RSS functionality to spread traffic across different XSK queues instead of having to provide explicit steering rules.	
TLS TIS Pool	[TLS-Enabled Devices] Per-connection hardware TIS objects is used to maintain the device TLS TX context. Use a SW TIS pool for recycling the TIS objects instead of destroying/creating them. This reduces the interaction with the device via the FW command interface, which increases the TLS connection rate.	
RDMA Featur	res	

Feature/Ch ange	Description	
Expand Rep Counters	[ConnectX-5 and above] Adding RDMA traffic-only counters for rep devices. These counters can now be read from host with ethtool or from sysfs and not only from the cointainer.	
UMR QP Recilency	[ConnectX-5 and above] Added a recovery flow for the driver's UMR logic so that other UMR requests can be proccessed after the error UMR was dropped and the UMR QP was reset. Previously, a faulty UMR request would have moved the QP to error state and disable any option to continue issuing UMRs.	
General	Bug fixes	

Feature/Ch ange	Description		
	5.8- 1.1.2.1		
General	Bug fixes		
5.8- 1.0.1.1			
Remove Dependen cy Between SR-IOV and eSwitch Mode	[All HCAs] Removed dependency between SR-IOV and eSwitch mode. Currently, there are three eSwitch modes: none, legacy, and switchdev (non of which are the default mode). When disabling SR-IOV, the current eSwitch mode will be changed to none. This feature removes eSwitch mode none and also removes dependency between SR-IOV and eSwitch mode.		
DevLink Parallel Command	[All HCAs] Added support for running DevLink commands in parallel on different DevLink devices is possible. For example, burning firmware on a few cards on the same host in parallel using DevLink API is now possible.		
Graceful Shutdown of Parent and Page Supplier	[All HCAs] Set default graceful period values for functions based on their type. ECPFs will get graceful period of 3 minutes, PFs get 1 minute, and VFs/SFs get 30 seconds.		
N Pulses Per Second (NPPS)	[ConnectX-6 Dx and above] Enhanced NPPS to allow setting a pulse period higher than 1 pulse per second and to allow setting the pulse width. If the width is unset, the driver implicitly sets it to half the given period (the width should be less than the pulse period). In this release, the pulse duration ranges between 65536 NS–524288 NS.		

Feature/Ch ange	Description	
Remote Invalidate Option for MKeys	[All HCAs] Addded support for the option to enable remote invalidation when creating a new mkey. This way the rkey for a memory region can be changed frequently.	
GPUDirect Over DMA-	 [All HCAs] Added support for GPUDirect support over dma-buf. As such, using the new mechanism nv_peer_mem is no longer required. The following is required for dma-buf support: Linux kernel version 5.12 or later OpenDM version 5.15 or later 	
BUF	 OpenRM version 515 or later Perftest support was added as well: Default option in perftest is without dmabuf. To run with this option, add use_cuda_dmabuf in addition to use_cuda flag. 	
Floating LID	[ConnectX-7] Added support for Floating LID (FLID) which can be used to identify a group of InfiniBand routers that allow communication with another subnet's entity. With this feature, multiple routers can be used per destination so that adaptive routing is supported. The FLID feature needs support from components such as the host, the subnet manager, the router, and more. This feature is only supported on the host portion of the sysrem.	
General	Bug fixes	

Feature/C hange	Description		
5.7-1.0.2.0	5.7-1.0.2.0		
Support Represen tor Metering Over SFs	[ConnectX-6 Dx and above and BlueField-2] Extended the support of representor metering from supporting only VFs representor to also supporting SFs representor.		
Exposing Error Counters on a VPort Manager	[ConnectX-4 and above] Added support for exposing error counters on a VPort manager function for all other VPorts. These counters can be used to detect malicious users who are exploiting flows that can slow the device. The counters are exposed through debugfs under: /sys/kernel/debug/mlx5/esw/ <func>/vnic_diag/</func>		

Feature/C	Description	
hange	Description	
Ethtool RX Flow Steering for IPolB	[All HCAs] Enabled steering of IPoIB packets via Ethtool, in the same way it i done today for Ethernet packets.	
Memory Consump tion Minimizat ion	[ConnectX-4 and above] Added support for providing knobs which enable users to minimize memory consumption of mlx5 functions (PF/VF/SF).	
XDP Support for Uplink Represen tors	ConnectX-5 and ConnectX-6 Dx) Added XDP support for uplink representors n switchdev mode.	
Resiliency to tx_port_t s	[ConnectX-6 Dx and above] Added resiliency to the tx_port_ts feature. private-flag may be enabled via ethtool tx_port_ts which provides a more accurate time-stamp. In very rare cases, the said time-stamp was lost, leading to losing the synchronization altogether. This feature allows for fast recovery and allows to quickly regain synchronization.	
Database of Devlink Health Asserts	[ConnectX-4 and above] Health buffer now contains more debug information like the epoch time in sec of the error and the error's severity. The print to dmesg is done with the debug level corresponding to the error's severity. This allows the user to use dmesg attribute: dmesg —level to focus on different severity levels of firmware errors.	
Expose FEC Counters via Ethtool	[ConnectX-5 and above] Exposed the following FEC (forward error detection) counters: ETHTOOL_A_FEC_STAT_CORRECTED • fc_fec_corrected_blocks_laneX • rs_fec_corrected_blocks ETHTOOL_A_FEC_STAT_UNCORR • fc_fec_uncorrectable_blocks_laneX • rs_fec_uncorrectable_blocks ETHTOOL_A_FEC_STAT_CORR_BITS • phy_corrected_bits	

Feature/C hange	Description	
	Command: ethtool -I show-fec <ifc></ifc>	
Applicati on Device Queues	[ConnectX-4 and above] Added driver-level support for Application Device Queues. This feature allows partition defining over the RX/TX queues into groups and isolates traffic of different applications. This mainly improves predictability and tail latency.	
Reinjectio n of Packets Into Kernel	[All HCAs] Added support for a new software steering action, mlx5dv_dr_action_create_dest_root_table(). This action can be used to forward packets back into a level 0 table. As a table with level 0 is the kernel owned table, this will result in injecting packets to the kernel steering pipeline.	
DCT LAG[ConnectX-6 Dx and above] Added firmware support to allow explice selection based on steering and not QP affinity. Functionality:DCT LAG1. Use LAG Hash Mode for the HCA with two ports, if supported 2. Keep port affinity function in LAG Hash Mode if it supports by select flow table in non-SwitchDev mode.		
AES-XTS in RDMA	Added support for plaintext AES-XTS DEKs.	
General	Bug fixes	

Customer Affecting Changes

23.10-1.1.9.0		
Customer Affecting Change	Description	
Lightweight Local SFs	Following the addition of the Lightweight Local SFs feature in version 23.07, in order to configure the scalable-functions, follow the revised instructions as detailed in the <u>Step-by-Step Guide</u> . Note: "Step 2.9 - Set all SF specific device parameters" is now mandatory for local SFs.	
23.10-0.5.5.0		
Customer Affecting Change	Description	

23.10-1.1.9.0		
Debugfs Directory Path Change	The debugfs directory of each interface can now be found under: /sys/kernel/debug/mlx5//, and not directly under the root of the debugfs filesystem (/sys/debug/kernel).	
Deprecation of OFED Public Power PC Installation	Starting from this release, MLNX_OFED releases for Power PC are no longer available for download from the public Download Center web page. Instead, you can find it on the following page: <u>https://network.nvidia.com/support/firmware/ibm-systemp/</u> .	
Pre-notification: Deprecation of Older Operating Systems	 Starting from next release, MLNX_OFED releases will no longer support operating systems with kernels below v4.18. This includes the following systems: Rhel7.x Debian9.13 Sles12.x Xenserver7.1 	
Customer Affecting Change	Description	
23.07-0.5.1.2		
Creating a QKEY with an MSB Set	To allow non-privileged users to create a QKEY with an MSB set, a new module parameter was added. For details, please see " <u>QKEY</u> <u>Mitigation in the Kernel</u> " under <u>New Features</u> .	
23.07-0.5.0.0		
IRQ Naming	IRQ renaming is no longer done when bringing the interface up/down. The IRQ name is now constant and is not affected by the interface state.	
RPM Packages Verification Key	<u>RPM_GPG-KEY-Mellanox</u> (or its variants) is no longer the public key that verifies RPM packages of MLNX_OFED. Instead, the RPM_GPG-KEY-Mellanox file on the top-level directory of the ISO should be used.	
Hairpin sysfs Support	Hairpin sysfs support was restricted to physical and virtual functions only.	
mlx5_core node_guid Module Parameter	Removed a non-functional mlx5_core node_guid module parameter.	
OpenSM Init	Starting from this release, the opensm init service moves from init.d (/etc/init.d/opensmd start) to systemd (# service opensmd start).	

23.10-1.1.9.0		
Apt Signing Key	Starting from this release, the public key that signed the apt repository of MLNX_OFED is included in the ISO in a format that can be used directly by the apt for repository signatures verification.	
IPoIB ULP Mode Deprecation	Starting from this release, MLNX_OFED supports IPoIB enhanced mode only. The ability to switch back to ULP mode using ipoib_enhanced module parameter is not supported. For more information about the enhanced mode, please refer to the OFED user manual, example: <u>Enhanced IP over InfiniBand</u> .	
Pre-notification: Deprecation of OFED Public Power PC Installation	Starting from next release, MLNX_OFED releases for Power PC will no longer be available for download from the public Download Center web page.	

Customer Affecting Change	Description	
23.07-0.5.0.0		
Creating a QKEY with an MSB Set	To allow non-privileged users to create a QKEY with an MSB set, a new module parameter was added. For details, please see "QKEY Mitigation in the Kernel" under <u>New Features</u> .	
IRQ Naming	IRQ renaming is no longer done when bringing the interface up/down. The IRQ name is now constant and is not affected by the interface state.	
RPM Packages Verification Key	<u>RPM_GPG-KEY-Mellanox</u> (or its variants) is no longer the public key that verifies RPM packages of MLNX_OFED. Instead, the RPM_GPG-KEY-Mellanox file on the top-level directory of the ISO should be used.	
Hairpin sysfs Support	Hairpin sysfs support was restricted to physical and virtual functions only.	
mlx5_core node_guid Module Parameter	Removed a non-functional mlx5_core node_guid module parameter.	
OpenSM Init	Starting from this release, the opensm init service moves from init.d (/etc/init.d/opensmd start) to systemd (# service opensmd start).	
Apt Signing Key	Starting from this release, the public key that signed the apt repository of MLNX_OFED is included in the ISO in a format that	

Customer Affecting Change	Description
	can be used directly by the apt for repository signatures verification.
IPoIB ULP Mode Deprecation	Starting from this release, MLNX_OFED supports IPolB enhanced mode only. The ability to switch back to ULP mode using ipoib_enhanced module parameter is not supported. For more information about the enhanced mode, please refer to the OFED user manual, example: <u>Enhanced IP over InfiniBand</u> .
Pre-notification: Deprecation of OFED Public Power PC Installation	Starting from next release, MLNX_OFED releases for Power PC will no longer be available for download from the public Download Center web page.

Customer Affecting Change	Description
23.04-0.5.3.3	
Netdev Interface Configurati on is not Preserved During Reload/Res et/Recovery	As of OFED 23.04, during reset/reload/recovery flows, the netdev interface is destroyed and re-created (rather than just suspended). As a result, the netdev interface configuration is not preserved, and must be re-applied. The way to do this is to use proper network-scripts and/or udev rules files to configure network interface parameters. These are automatically triggered whenever a netdev interface is added, regardless of whether it was added due to a user-initiated operation or an automatic failure recovery operation. Thus, no special processing is required to re-apply the network interface configuration parameters following a reset/reload/recovery operation – it is performed automatically.
Prenotificati on: Deprecatio n of OFED Public Power PC Installation	Starting next release, MLNX_OFED releases for Power PC will no longer be available for download from the public Download Center web page.
Prenotificati on: ULP Mode Deprecatio n	Starting next release, MLNX_OFED will support IPoIB enhanced mode only. The ability to switch back to ULP mode using ipoib_enhanced module param will not be supported. For more information about the enhanced mode, please refer to OFED user manual, example: <u>Enhanced IP over InfiniBand</u>

Customer Affecting Change	Description
	In order to address RHEL kernel symbol changes, ISO images for the following operating systems are built with the updated kernel versions as follows:
	OS Name Old Kernel
	New Kernel
	rhel8.6-aarch64 4.18.0-372.9.1.el8_6.aarch64
	4.18.0-372.41.1.el8_6.aarch64
	rhel8.6-ppc64le 4.18.0-372.0.1.el8_6.ppc64le
	4.18.0-372.41.1.el8_6.ppc64le
	rhel8.6-x86_64 4.18.0-372.9.1.el8_6.x86_64
	4.18.0-372.41.1.el8_6.x86_64
	rhel8.7-aarch64 4.18.0-425.3.1.el8.aarch64
	4.18.0-425.14.1.el8_7.aarch64
	rhel8.7-ppc64le 4.18.0-425.3.1.el8.ppc64le
lastallation	4.18.0-425.14.1.el8_7.ppc64le
Installation, ISO, RedHat	rhel8.7-x86_64 4.18.0-425.3.1.el8.x86_64
	4.18.0-425.14.1.el8_7.x86_64
	rhel9.0-aarch64 5.14.0-70.13.1.el9_0.aarch64
	5.14.0-70.46.1.el9_0.aarch64
	rhel9.0-ppc64le 5.14.0-70.13.1.el9_0.ppc64le
	5.14.0-70.46.1.el9_0.ppc64le
	rhel9.0-x86_64 5.14.0-70.13.1.el9_0.x86_64
	5.14.0-70.46.1.el9_0.x86_64
	rhel9.1-aarch64 5.14.0-162.6.1.el9_1.aarch64
	5.14.0-162.19.1.el9_1.aarch64
	rhel9.1-ppc64le 5.14.0-162.6.1.el9_1.ppc64le
	5.14.0-162.19.1.el9_1.ppc64le
	rhel9.1-x86_64 5.14.0-162.6.1.el9_1.x86_64
	5.14.0-162.19.1.el9_1.x86_64
	This change comes to support DedHat undated kernels without the need
	This change comes to support RedHat updated kernels without the need to addadd-kernel-support during OFED installation.

Customer Affecting Change	Description
Power Setups on UCX/HPC-X	UCX/HPC-X no longer supports Power setups.
NEO-Host	Starting from this release, OFED will discontinue the provision of NEO- Host. NEO-Host can be manually downloaded and installed using the following guide: <u>https://docs.nvidia.com/networking/display/NEOSDKv26/Installation+and+I</u> <u>nitial+Configuration#InstallationandInitialConfiguration-</u> <u>DownloadingtheMellanoxNEOSDKSoftware</u>
dapl	Starting from this release, OFED will discontinue the provision of dapl.
Signing Key for SLES15 sp4 and sp5	As of version 23.04, the builds for SLES15 sp4 and sp5 are being singed with a newer signing key. The corresponding public key can be downloaded from https://www.mellanox.com/downloads/ofed/nv_nbu_kernel_signing_key_pu b.der instead of https://www.mellanox.com/downloads/ofed/mlnx_signing_key_pub.der.
dump_pr SM Plugin	Starting from this release, OFED will discontinue the provision of dump_pr subnet manager plugin.
mpi- selector	Starting from this release, OFED will discontinue the provision of mpi-selector.
OpenSM Init	Starting 23.07 release, opensm init service will move from init.d to systemd.

Custom er Affectin g Change	Description		
5.9-0.5.6	5.9-0.5.6.0		
Deprec ation, LAG Mode	Setting LAG mode via Sysfs is going to be deprecated in a future release. Instead, LAG Hash mode will be used by default, similar to upstream behavior.		

Custom er Affectin g Change	Description
via Sysfs	
LAG Configu ration, PCI Error	From version 5.9, LAG configuration will be lost in case driver incurs a PCI error. Make sure to reconfigure the bond after driver completes the recovery from the PCI error. In releases prior to 5.9, in case of PCI error (EEH injections on PPC setup), the driver recovers LAG bond and reconfigures it automatically in case it what configured before the appearance of the error.

Custom er Affecti ng Change	Description		
5.7-1.0.2	5.7-1.0.2.0		
Multi- Block Encrypt ion	Multi-block encryption is currently unsupported, due to a hardware limitation.		

Feature /Chang e	Description	
5.6-2.0.9	5.6-2.0.9.0	
Operati ng System s	Added support for the following Operating Systems: RHEL8.6, RHEL9.0, SLES15-SP4.	
General	Bug fixes	

Feature /Chang e	Description	
5.6-1.0.3	.3	
General		
New Adapter Card Suppor t	Added support for ConnectX-7 adapter cards. ConnectX-7 has the same feature set as the ConnectX-6 adapter card.	
ASAP ² F	eatures	
Bridge Spoof Check	[All HCAs] Added support for spoof check with TC flower rules on representors attached to bridge to mirror spoof check SR-IOV functionality.	
Setting VF Group Rate Limit	[ConnectX-5 and above] Added support for setting VF group rate limit using Devlink command.	
TC Flows on Shared Block	[ConnectX-5 and above] Added support for creation of TC flows on shared block of VF representors.	
Flow Meterin g	 [ConnectX-6 Dx and above] Added support for offloading OpenFlow Meters in OVS-DPDK. Please note the following: Meter offload can be applied only on port 0 and it's VFs Only one meter per flow is allowed Only one meter band per meter is allowed Only meter band type drop is supported Meter-stats might not be accurate 	
Core Fea	Core Features	
Firmwar e Reset	[BlueField-2] Added support of firmware reset in DPU NIC mode.	
Increas ed	[All HCAs] Increased the firmware pre-initialization timeout from 2 minutes to 2 hours when waiting for firmware during driver health recovery, allowing the	

Feature /Chang e	Description
Robust ness of mlx5_c ore Driver Recover y	driver to passively recover from a firmware reset, even if the reset takes an unusually long time. Additionally, added an exit clause to the wait for firmware loop, allowing immediate response to a user initiated device removal.
NetDev F	Features
Ethtool CQE Mode Control	[ConnectX-4 and above] Replaced the vendor-specific Ethtool API (priv-flag) with a standard Ethtool API (replaced 'ethtoolset-priv-flags ethX rx_cqe_moder on/off tx_cqe_moder on/off' with 'ethtool -C ethX cqe-mode-rx on/off cqe-mode-tx on/off'). This decreases the amount of vendor-specific configurations and aligns mlx5 driver with the upstream Ethtool API.
SyncE	[ConnectX-6 Dx] Added an indication in SyncE Daemon that states whether SyncE engine moved to holdover state due to failure (the reason for failure will be displayed). In addition, added indication whether SyncE engines collected enough frequency samples in order to move to holdover. Note: Not supported in ConnectX-6 Lx adapter cards.
RDMA Fe	eatures
VFIO, CQ Interru pt Mode	[ConnectX-5 and above] Added support for VFIO applications to listen on and capture completion events via the Event Queue mechanism.
VFIO, Asynch ronous Event	[ConnectX-5 and above] Added support for VFIO applications to listen on and capture device asynchronous events via the Event Queue mechanism.
Security	
OVS- IPSec Full Offload	[BlueField-2] Added support for configuration of IPsec full offload using OVS by adding VXLAN tunnel to OVS with the PSK option.
Software	e Steering Features
Full Tunnel	[ConnectX-6 Dx and above] Added support for using full-tunnel-header matching along with many other criteria within one matcher. This feature uses

Feature /Chang e	Description
Header Matchi ng	the new definer index, defined in the firmware, to build a matcher so that the full tunnel header matching can be used along with all other criteria.
Matchi ng Granula rity Change	[ConnectX-5 and above] Added support for matching granularity change. As a result, when creating FDB flow with destination of VPORT, a src_port matching must no longer be added. Now, FDB flow can match all vports and goto a VPORT destination. The new behavior is the same as done on firmware steering.
Installati	on
Installat ion	New options were added to the ofed_uninstall.sh script:only-kernel andonly- user. Those can be used to uninstall only kernel packages or only user-space packages (the equivalent of kernel-only install or user-space-only install, respectively). This may be useful to keep different sets of kernel and user- space installations. When running the uninstall script with a combination ofonly-kernel andonly- user produced an undefined result.

Custom er Affectin g Change	Description
Custome	r Affecting Changes 5.6-1.0.3.3
Interfac e Renami ng, PF/VF, Udev	The OFED driver no longer performs Ethernet NetDev interface renaming for PFs and VFs. The udev rules file which implemented renaming (82-net-setup-link.rules) and its supporting script vf-net-link-name.sh are no longer installed by default. Renaming is thus performed by underlying mechanisms in udev, in the kernel, and in the BIOS. Users who wish to continue using the OFED driver renaming mechanism must add optioncopy-ifnames-udev to the OFED install command. To install these files at a later time, copy them from one of the following directories: • /usr/share/doc/mlnx-ofa_kernel (RHEL8 and newer) • /usr/share/doc/mlnx-ofa_kernel-[1-9]* (RHEL 7.X) • /usr/share/doc/packages/mlnx-ofa_kernel (SLES) • /usr/share/doc/mlnx-ofed-kernel-utils/examples (Debian-based releases)

Custom er Affectin g Change	Description	
	 Note File 82-net-setup-link.rules should be copied to directory /etc/udev/rules.d File vf-net-link-name.sh should be copied to directory /etc/infiniband (make sure that it has both read and execute permission) After copying over the files, the driver should be restarted for the copied files to take effect Customers who wish prevent renaming of NetDev names should add "net.ifnames=0 biosdevname=0" to the kernel boot command line, and then reboot the host 	
Commu nity Operati ng System s	Starting OFED 5.6, NVIDIA is introducing a new support model for OFED used on open source community operating systems. The goal of this new support model is to enable customers to use community-maintained variants of the Linux operating system, without being limited to major distributions that NVIDIA provides primary support for. For more information, see " <u>Installation on Community Operating Systems</u> " section in the user manual. For a list of supported Community OSs, please see " <u>Supported Community Operating</u> <u>Systems</u> " section in the release notes.	
ar_mgr Subnet Manage r Plugin	MUNX (JEEL) USER MANUAL	
Fabric Collect or in UFM	Starting UFM v6.7, Fabric Collector is no longer supported. For more information, see the <u>UFM release notes</u> .	
	(j) Note	

Custom er Affectin g Change	Description
	Please note that UFM is no longer bundled with OFED.
OVS- DPDK— Partial Offload	Starting OFED 5.6, OVS-DPDK does not support partial offload.

Customer Affecting Change	Description
5.5-1.0.3.2	
Disabling RoCE While Using sysfs	When using sysfs to enable/disable roce in kernel 5.5 and up, the "devlink reload" command (using iproute2 with devlink tool) will need to be used to activate the RoCE status change. Disable RoCE example: 1. echo 0 > /sys/bus/pci/devices/0000:08:00.0/roce_enable 2. devlink dev reload pci/0000:08:00.0
mlnx- ofa_kernel Installation	The source code for mlnx-ofa_kernel is no longer installed by default on RPM-based distributions (e.g., RHEL and SLES). Notes: • mlnx-ofa_kernel is included in the < <package mlnx-ofa_kernel-devel-<br="">source>> in the MLNX_OFED distributions under RPMS/ and may be manually installed from there. • There is no change for deb-based distributions (Debian and Ubuntu). The full source is included, as before, in the package mlnx-ofed-kernel- dkms.</package>
Software Encapsulation Compatibility	There is an encapL2 compatibility issue with accelerated reformat action creation using mlx5dv_dr API. Using OFED 5.4 with firmware xx.32.1xxx and above or using OFED 5.5 with firmware lower than xx.32.1xxx will not allow accelerated reformat action. (Using OFED 5.4 and 5.5 with bundle firmware works properly.)
xpmem in RHEL8	Added xpmem packages in RHEL8 builds.

Customer Affecting Change	Description
Python3	Starting OVS DPDK 2.15, the Python minimum required version is 3 and OVS-DPDK will not be compiled using Python 2.

Customer Affecting Change	Description	
Customer Affecting Changes 5.4-3.0.3.0		
CUDA, UCX, HCOLL	For UCX-CUDA and hcoll-cuda, CUDA was upgraded from version 10.2 to 11.2.	

Customer Affecting Change	Description
Customer Affecting	Changes 5.4-1.0.3.0
udev Rules	As of version 5.4, the driver is set so that udev rules will change the names of network interfaces created from NVIDIA adapters. The udev rules are shipped to "/lib/udev/rules.d" and may be overridden by placing a file with the same name in "/etc/udev/rules.d". Example: /etc/udev/rules.d/82-net-setup-link.rules
Network Interface Names, udevd	[ConnectX-4 and above] In MLNX_OFED 5.4 GA, ConnectX-4/5/6 Ethernet network interfaces are now provided with permanent names. Prior to this release, the default interface names were provided by the kernel and udevd (ethX) remained as-is. From this release onwards, interface names are generated via new udevd rules. The generated names are now predictable, and the default names are automatically renamed to the predictable names by the udevd daemon, according to udev rules files installed by OFED. The new interface names look as follows: en[P][p <bus number]sf<br="">For example, a ConnectX device with PCI address: 0005:01:00.1 will be named enP5p1s0f1 The advantage of such a scheme for interface naming is that device whose PCI address is 0005:01:00.1 will always get the same device name since that name now depends on the host geography. (Previously there were race conditions which sometimes caused the same physical device to get a different interface name upon reboot).</bus>

Customer Affecting Change	Description
	Note: Ethernet interface renaming for ConnectX-4/5/6 is performed only when eswitch is supported. eswitch is supported on kernels starting from kernel version 4.9; for Linux distro kernels earlier than 4.9.0, eswitch is supported only on RHEL7.x and on XenServer 7.1 CU2.
Deprecated, OvS- DPDK	OvS-DPDK deprecated the command "ovs-appctl dpctl/dump-e2e- stats". Instead, the command has been integrated into the existing command "ovs-appctl dpctl/offload-stats-show -m" (when e2e- cache is enabled).
OvS-DPDK	OvS-DPDK ct-ct-nat offloads is now disabled by default. A new knob in OvS was introduced: "ovs-vsctl set open_vswitch . other_config:ct-action-on-nat-conns=" (default value is false). If disabled, ct-ct-nat configurations will not be fully offloaded, improving connection offloading rate for other cases (ct and ct- nat). If enabled, ct-ct-nat configurations will be fully offloaded but ct and ct-nat offloading will be created more slowly.
mlnxofedinstall, udev, MLNX_OFED, umad	Before version 5.4, /etc/udev/rules.d/90-ib.rules was potentially automatically edited by installation scripts in case the options umad-dev-rw orumad-dev-na were used. From version 5.4 and above, those changes are made in /etc/udev/rules.d/91-ib- permissions.rules which (if exist) only include the settings for those command-line options.

Feature/Change	Description	
5.5-1.0.3.2		
ASAP ² Features		
Bridge Offloads with VLAN	[ConnectX-4 and above] Added support for bridge offloads with VLAN support that works on top of mlx5 representors in switchdev mode.	
Supporting OVS Groups in Fast- Failover Mode	[ConnectX-6 Dx] Improved OVS failover through support for OVS groups in fast-failover mode + VF_LAG configuration with OVS.	
Exposing Hairpin Queues Information	[ConnectX-6 Dx and BlueField-2] Added support for exposing hairpin out of buffer drop counter per device. This feature shows buffer drops related only to hairpin queues which were opened on the queried device.	

Feature/Change	Description
	To enable this counting mode (this must be done before any hairpin rules are created), use the following: echo "on <peer_devname>" > /sys/class/net/<dev>/hp_oob_cnt_mode where <peer_devname> is the peer device to which traffic coming to the configured device will be forwarded to for transmission. To read the drop counter, use the following: cat /sys/class/net/<dev>/hp_oob_cnt</dev></peer_devname></dev></peer_devname>
Linux Bridge Offload	[ConnectX-6 Dx and BlueField-2] Added bridge offloads to support bonding (VF LAG), attaching bond device to bridge instead of uplink representors.
VLAN Pop/Push	[ConnectX-6 Dx] Added OOB support for VLAN push on Rx (wire to VF) and VLAN pop on Tx (wire to VF) in switchdev mode.
Offload Forwarding to Multiple Destinations	[ConnectX-5 and above] Added support for offloading packet replication to up to 32 destination through the use of TC rule.
Slow Path Metering	[ConnectX-4 and above] Expanding the RDMA statistic tool to support setting vendor-specific optional counters dynamically using netlink. Added to mlx5_ib the following optional counters: cc_rx_ce_pkts,cc_rx_cnp_pkts,cc_tx_cnp_pkts. Example: \$ rdma statistic mode supported link rocep8s0f0/1 link rocep8s0f0/1 supported optional-counters cc_rx_ce_pkts,cc_rx_cnp_pkts,cc_tx_cnp_pkts \$ sudo rdma statistic set link rocep8s0f0/1 optional-counters cc_rx_ce_pkts,cc_rx_cnp_pkts \$ rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts,cc_rx_cnp_pkts \$ sudo rdma statistic set link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts,cc_rx_cnp_pkts \$ sudo rdma statistic set link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts,cc_rx_s \$ sudo rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts \$ rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts \$ sudo rdma statistic mode link rocep8s0f0/1 link rocep8s0f0/1 optional-counters cc_rx_ce_pkts \$ sudo rdma statistic unset link rocep8s0f0/1 optional-counters
Using Specific Interface for Tunnel Route Lookup	[ConnectX-5 and above] Added ability to use a specific interface for tunnel route lookup if tunnel was created with the current device.
Core Features	

Feature/Change	Description
Subfunction Trust Configuration Enhancement	[ConnectX-5 and above] Added support via mlxdevm to mark a given PCI subfunction (SF) or virtual function (VF) as a trusted function. The device/firmware decides how to define privileges and access to resources.
Prevent VF Memory Exhaustion	[All] Added support for preventing VF memory exhaustion. This feature exposes a sysfs (to the system admin) which can set a limit on each VF memory consumption. Note: Currently only supported on Ethernet.
BlueField NIC Separate Reset	[BlueField-2] Added support for resetting the NIC domain of BlueField-2 while keeping ARM alive.
Multiple Steering Priorities for FDB Rules	[ConnectX-6 Dx and BlueField-2] Added support in multiple flow steering priorities for FDB rules.
NetDev Features	
Traffic Engineering: Hierarchical QoS	[ConnectX-5 and above] Added support for offloading the HTB qdisc to the NIC, allowing it to scale better by eliminating a single locking point. The configuration is done with the TC commands. Note: Kernel 5.15 or higher is required. Limited to 256 nodes.
TLS RX Resynchronizati on Resiliency Feature Description	[ConnectX-6 Dx and above] Added support for driver resiliency against high load of RX resync operations.
Simultaneous PTP and CQE Compression	Added support for the activation of PTP and CQE compression simultaneously. Since CQE compression might harm the accuracy of the PTP, the feature enables PTP packets to be moved to a dedicated queue where they are not subjected to compression. However, this configuration conflicts with setting aRFS. Turning off CQE compression, causes a hiccup in traffic which may cause a loss of synchronization. To overcome this, restart the synchronization. Note: This combination is supported only for Ethernet drivers. Other driver profiles, like IPoIB and representors, do not support this combination.
RDMA Features	
ODP On Demand Synchronization	Added support to expose an option to prefetch ODP MR without faulting. This enables updating the device page table with the presenting CPU pages and reducing page faults in the system.

Description
[ConnectX-6 and above] Added DV API that allows configuration of MKey with AES-XTS crypto offloads. The MKey can be configured for both crypto and signature offloads.
[ConnectX-4 and above] Added support to allow DEVX UMEM to be created with larger page sizes than 4K. For some device objects (e.g., RegEx) this is a must. In addition, page size larger than 4K may need less MTTs which may improve performance. A new API mlx5dv_devx_umem_reg_ex() was added which requests a specific page sizes. It enables better application control on the required UMEM page size. The new API named mlx5dv_devx_umem_reg_ex() will be part of rdma-core V35.
[ConnectX-4 and above] Added support for cleanup of the synchronize_srcu() from the ODP flow because it was a time- consuming part of dereg_mr. Note: This only affects the driver and not the firmware.
[ConnectX-4 and above] Extended support for the "rdma res show" command to SRQ and context resources.
[ConnectX-5 and above] Added support for Raw WQE (mlx5dv_wr_raw_wqe). This feature allows applications to build a new custom work request (WQE) that is not supported by the verbs or driver and post it on normal QP. It is an extension for IBV work request (ibv_wr_*) with mlx5 specific features for sending a work request.
Added support for mlx5 user space driver over VFIO. This feature enables an application to take full ownership on the opened device and run any firmware command (e.g., port up/down) without any concern to hurt someone else. The application look and feel is like regular RDMA application over DEVX. It uses verbs API to open/close a device and then mostly uses DEVX APIs to interact with the device. New mlx5 DV APIs were added to get ibv_device for a given mlx5 PCI name and to manage device specific events. For description of the relevant APIs and expected usage of those APIs, look up the following: mlx5dv_get_vfio_device_list() mlx5dv_vfio_get_events_fd() mlx5dv_vfio_process_events()
<u>I</u>

Feature/Change	Description		
Bonding Interfaces	With some specific NICs, each interface may have different PCIe domain, bus, device, or function IDs. For interfaces with the same system_image_guid, the driver assumes they reside on the same physical device and use a native_port_id to distinguish its index. Fallback is to PCIe BDF, if unsupported.		
Software Steering Dump File Parser Tool	[ConnectX-4 and above] mlx_steering_dump tool is used to parse software steering dump files which includes information about domains, tables, matchers and rules created by software steering (mlx5dv_dr API), it can be used offline by providing a dump file as input, or it can be used to trigger DPDK app (like testpmd) to generate the dump and parse it.		
Installation Featu	Installation Features		
Multiple Development Headers Packages	Allowed installing multiple mlnx-ofa_kernel development headers packages (for different kernel versions of the same mlnx-ofa_kernel package version) side by side on the same system.		
Kernel Module Signature	Added signature of kernel modules of EulerOS 2.0 SP8-SP10 (x86_64 and aarch64) builds of MLNX_OFED.		
Enable sf-cfg- drv by Default in EulerOS2.0	Enabled SF_CFG (SF config dummy driver,with-sf-cfg-drv) on EulerOS2.0 SP8 and SP10.		

Feature/Change	Description		
5.4-3.0.3.0	5.4-3.0.3.0		
GPUDirect			
Kernel Space GPUDirect from VM	[ConnectX-5 and above] Added support for kernel space GPUDirect from the VM. To use GDS with high performance in a VM, set the ATS capability in ib_alloc_mr.		
ASAP ²			
CTC Metering	[ConnectX-6 Dx] Added support for per flow metering using OVS or TC, PPS, and BPS.		
Slow Path Metering	[ConnectX-6 Dx] Added support for slow path metering on representor using OVS or TC, PPS, and BPS.		
Core			

Feature/Change	Description
VF LAG	[ConnectX-6 Dx and BlueField-2] Added support to have physical port selection based on the hash function defined by the bond so that different packets of the same flow will be egress from the same physical port. In order to enable this feature, set this mode for both bonded devices through the below sysfs before the device is in switchdev mode: echo "hash" > /sys/class/net/enp8s0f0/compat/devlink/lag_port_select_mode In order to have the legacy behaviour (queue affinity based selection), echo the following: echo "queue_affinity" > /sys/class/net/enp8s0f0/compat/devlink/lag_port_select_mode This feature requires to set LAG_RESOURCE_ALLOCATION to 1 with mlxconfig.
Single IRQ for PCI Function	[ConnectX-4 and above] Added support for single IRQ for PCI function. To use a high number of VFs, a large amount of IRQs is required which the device cannot always support. This feature enables VFs to function with a minimum of a single IRQ instead of two. This is done via dynamic MSIX feature. In case dynamic MSIX feature is not supported (old kernels), the following configuration will probe all VFs with single IRQ: \$ mlxconfig -d <pci_dev> s NUM_VF_MSIX=0 STRICT_VF_MSIX_NUM=1</pci_dev>
Netdev	
ethtool EEPROM Support for DSFP	[ConnectX-6 Dx] Added support for reading DSFP module information. The change includes adding new options to ethtool netlink EEPROM module read API to read a specific page and bank.
RDMA	
Dynamic VF MSI-X Allocation	[ConnectX-5 and above] Added support for dynamic assignment of MSI-X vector count. The number of MSI-X vectors is a PCI property visible through lspci and is read-only field configured by the device. The static assignment of an amount of MSI-X vectors does not allow to utilize the newly created VF because the future load and configuration where that VF will be used is not known to the device. The VFs are created on the hypervisor and forwarded to the VMs that have different properties (for example number of CPUs). To overcome the inefficiency in the spread of such MSI-X vectors, the kernel is now allowed to instruct the device with the needed number of such vectors before the VF is initialized and bounded to the driver.
DV API for DMA GGA memcpy	[BlueField-2 and above] DMA memcpy is one of several Memory-to- Memory Offloads (MMO) available from BlueField-2 onwards. It utilizes

Feature/Change	Description
	the GGA modules on the DPU to perform DMA memcpy, thus improving performance. The memcpy can be done locally, on the same host, or between the host and the Arm. To use this feature, expose DV API.
Steering UserSpa	ace
Set DR Matcher Layout	[ConnectX-6 Dx] Added support for a new RDMA CORE DR API to set the DR matcher layout by calling mlx5dv_dr_matcher_set_layout. Setting the matcher layout allows presetting the matcher size and increasing matcher rule capacity, as well as other performance improvements in case matcher size is known.
Flex Parsers misc4	[ConnectX-5 and ConnectX-6 Dx] Added ability to expose flex parsers 4-7 provided by misc4 to extend matching ability of flex parsers. Now all flex parsers can be matched at the same time.
Software encap Action	[ConnectX-5 and above] Added support for software encap action. There is requirement for more than 1M encap actions, but currently the encap action creation uses devx, which is very slow for 1M encap actions. As such, there is a need to support a way for software to create encap actions. The encap reformat action creation in rdma-core can now be done via software, rather than devx. It will use the new ICM memory type of software encap and directly copy encap data there, then use the memory pointer for flow creation.
Bug Fixes	See <u>Bug Fixes</u> .

Feature/Cha nge	Description
5.4-1.0.3.0	
ASAP ²	
Enlarge Switchdev Tables	[ConnectX-5 and above] Added support for allowing OVS kernel to support up to 128 matches (groups) per table and 16M entries per group.

Feature/Cha nge	Description
Offloading Extended ct_state Flags	 [ConnectX-5 and above] Added support to offload ct_state flags rpl, inv, and rel. For rpl, support was added for both set and not set matching offload (i.e., +rpl and -rpl). For inv and rel, support was added only for the not set option (i.e., -rel and -inv).
Core	
Scalable Functions (Subfunctio ns)	[ConnectX-5 and above] Added support for scalable functions (also called subfunctions). The feature enables the user to create, configure, and deploy a scalable functions (e.g., RDMA and networking applications) and to assign them to a container when a container is started via mlxdevm tool. A scalable function can also be deployed in an untrusted guest/host system from the NIC/DPU. This enables full configuration of the function and its representors from the NIC/DPU before giving the function for a container to run in a host system. For more information, see <u>https://github.com/Mellanox/scalablefunctions/wiki/MLNX_OFED-step-by- step-guide</u> .
Scalable Function QoS	[ConnectX-5 and above] Added support for scalable function QoS and QoS group via mlxdevm's rate commands. Run "man mlxdevm port" for details.
Auxiliary Bus in mlx5 Driver	[ConnectX-4 and above] Updated mlx5 driver to use auxiliary bus in order to integrate different driver components into driver core and optimize module load/unload sequences.
Installation	
Script Removal from mlnx- ofa_kernel	[General] Moved all Python scripts and some other common scripts out of the mlnx-ofa_kernel packages. This removed the python dependency from that package when rebuilding it and avoided unnecessary errors when rebuilding them for custom kernels.
Netdev	
What-Just- Happened (WJH) in NICs	[ConnectX-4 and above] Added support for WJH in NICs. WJH allows for visibility of dropped packets (i.e., receiving notice of drop counters increase, seeing content of the dropped packets, debugging, and more). WJH is a service in devlink context and it is already implemented in the switch.

Feature/Cha nge	Description
	Note: processing dropped packets (even for visibility purposes) may cause a degradation in performance and leaves the driver vulnerable for malicious attacks. The feature is disabled by default. Supported traps:
	 VLAN mismatch: existing generic trap DEVLINK_TRAP_GENERIC_ID_DMAC_MISMATCH Traps received packets with wrong VLAN tag DMAC mismatch: new generic trap DEVLINK_TRAP_GENERIC_ID_DMAC_MISMATCH Traps received packets with wrong destination MAC
	Support added in user-space (N/A or package name + version): Devlink infrastructure (man7.org/linux/man-pages/man8/devlink-trap.8.html) Devlink provides an infrastructure called devlink trap which allow a device to register/unregister and to enable/disable traps. Devlink traps also provide traps grouping and policing. The trapped packets are monitored and then forward to the drop monitor. Drop monitor is used to send notifications to user space about dropped packets. Note: For this release, NIC WJH will not implement the policy.
	[General] Added ethtool extended link state to mlx5e. ethtool can be used to get more information to help troubleshoot the state. For example, if there is no link due to missing cable, run the following: \$ ethtool eth1
ethtool Extended Link State	 Link detected: no (No cable) Besides the general extended state, drivers can pass additional information about the link state using the sub-state field. Example: \$ ethtool eth1
	 Link detected: no (Autoneg, No partner detected) The extended state is available only for some cases of no link. In other cases, ethtool will print only "Link detected: no" as it did before.
RDMA	
DV "Signature API"	[ConnectX-5 and above] Added support for "Signature API" which, on supported devices, allows application-level data-integrity checks via a signature handover mechanism. Various signature types, including CRC32 and T10-DIF, can be automatically calculated and checked, stripped, or appended during the transfer at full wire speed.

Feature/Cha	
nge	Description
ibv_query_q p_data_in_o rder() verb	[General] Added support for ibv_query_qp_data_in_order() API. This API enables an application to check if the given QP data is guaranteed to be in order, enabling poll for data instead of poll for completion.
Relaxed Ordering for Kernel ULPs	[ConnectX-4] Added support for enabling Relaxed Ordering for Kernel ULPs. Using relaxed ordering can improve performance in some setups. Since kernel ULPs are expected to support RO, it is enabled for them by default so they can benefit from it.
ah_to_qp Mapping	[ConnectX-6 Dx] Added support for mapping a QP to AH over DEVX API, which enables DC/UD QPs to use multiple CC algorithms in the same data center.
Steering Use	rSpace
Matching on RAW Tunnel Headers	[ConnectX-5 and above] Added DR support for matching on RAW tunnel headers using the misc5 parameters, This feature allows matching on each bit of the header, inducing reserved fields.
Software Steering Insertion Rate Optimizatio ns	[ConnectX-6 Dx] Added support for better insertion rate in software steering. This includes multi-QP which skips areas in the code that may be for debug only.
Software Steering Rule Optimizatio n	[ConnectX-6 Dx] Improved rate of updating steering rules, insertion, and deletion. The feature includes definers, multi-qp approach, and better memory usage.
Duplicate Rules Insertion	 [ConnectX-5 and above] Added support for ability to allow or prevent insertion of duplicate rules, so the user can choose one of the following behaviors: 1. Prevent duplicate rules, so that already-existing rule and fail can be detected. 2. Allow duplicate rules, to enable updating the rule's action (this will only take effect once the previous rule is deleted). By default, duplicate rules are allowed.
Improved Software Steering Rule	[ConnectX-6 Dx] Made it so that all rule's insertion occur in a defined time using defined (export) size of Htble and decreased use of dynamic allocation.

Feature/Cha nge	Description
Creation Stability	

Feature/Cha nge	Description
5.3-1.0.0.1	
MLX5DR SF	[ConnectX-5 and above] Added support for up to 512 SFs with the mlx5dv_dr API.
Dump Single Flow	[ConnectX-5 & ConnectX-6 Dx] Added support to dump single flow/rule with flow-id.
Local/Remo te Mirroring	[ConnectX-5* & ConnectX-6 Dx/ BlueField* & BlueField-2] OVS-DPDK added support for local and remote mirroring for offloaded traffic. *Enabling the port mirroring feature on a ConnectX-5 NIC and BlueField will break Connection Tracking.
Connection Tracking Replay State	[BlueField-2] Added support for matching on CT state replay.
Kernel TLS Offload	[BlueField-2] Added support for TX and RX kTLS offloads on the ARM in switchdev mode via a sub-function.
Increased Number of Virtio Functions	[ConnectX-6 Dx & BlueField-2] Added support for up to 504 Virtio functions. 512 total functions are supported, but some are consumed by PF, Host PF, and RSHIM.
VF Metering	[ConnectX-6 Dx] Added support for RX/TX metering per VF using sysfs API.
PTP Hardware Translation Offload	[ConnectX-6 Dx] Added support for the hardware clock device to be adjusted and provide timestamps which are translated into real-time nanoseconds. This can be used by the driver for PTP protocol. For further information, see <u>PTP Cyc2time Hardware Translation Offload</u> section.
TLS Rx Hardware Offload	[ConnectX-6 Dx] Added GA-level support for hardware offload decryption of TLS Rx traffic over crypto-enabled ConnectX-6 Dx NICs and above. Note: Not supported in ConnectX-6 Lx adapter cards.
MLX5DR Match	[ConnectX-6 Dx] Added support for match definers which are used internally in the mlx5dv_dr API. Definers allow filtering on more packet

Feature/Cha nge	Description
Definer	fields, improving the packet rate and accelerates mlx5dv_dr API.
MLX5DR Packet OK and Checksum Checks	[ConnectX-6 Dx] Added support for new matching fields ipv4_checksum_ok and I4_checksum_ok. I3_ok, I4_ok.
Pop VLAN on VF/SF Tx Direction	[ConnectX-6 Dx] Added support to pop VLAN on VF/SF Tx direction.
Connection Tracking Window Validation	[ConnectX-6 Dx, ConnectX-6 Lx, BlueFlield-2] Added support for ASO connection tracking of action creation and modification. This action allows performing TCP connection tracking using hardware offloads. Using this offload, the validity of the connection state of the incoming or outgoing packets on this TCP connection can be examined. Also added the ability for an ASO CT action created on one GVMI to be used on different GVMI.
Pyverbs	[All HCAs] Pyverbs are no longer being built for Debian 9.
Bug Fixes	See <u>Bug Fixes</u> .

Category	Description	
5.2-2.2.0.0		
DCT Support for Connection Establishm ent with RDMA_CM	[ConnectX-5 and above] [Alpha level] Added support for the dv APIs to allocate/deallocate a unique QP number that can be used as DC QPN in RDMA_CM connection establishment.	
МХМ	[All HCAs] The MXM package is deprecated and removed from MLNX_OFED.	
Devlink Firmware Reset	[All HCAs] In MLNX_OFED v5.2-1.0.4.0, it was noted that MLNX_OFED did not include the latest iproute2 that provided support for this feature, and that the latest iproute2 must be installed from Github (see <u>Changes and</u> <u>New Features History</u> section). This note is no longer relevant.	
Bug Fixes	See <u>Bug Fixes</u> .	
5.2-1.0.4.0		

Category	Description	
Rx Multi- strides CQE Compressi on	[ConnectX-5 and above] Added CQE compression support for Rx multi- strides packets.	
Multi- application QoS	[ConnectX-5 and above] Added support for configuring QoS on a single QP or on a group of QPs.	
MPLS-over- UDP Hardware Offload Support	[ConnectX-5 and above] Added support for encap/decap hardware offload of IPv4 traffic over MPLS-over-UDP. This can be used in networks with MPLS routers to achieve more efficient routing.	
Connection Tracking with Hairpin	[ConnectX-5 and above] Added support for adding connection tracking rules on VFs to forward traffic from one VF to the other.	
sFlow Sampling Rules Offload	[ConnectX-5 and above] Added support for offloading sFlow sampling rules. sFlow is an industry standard technology for monitoring high speed switched networks.Open vSwitch integrated sFlow to extend the visibility into virtual servers, ensuring data center visibility and control. Added support for offloading sFlow sampling rules.	
mlx5dv_dr Software Steering Parallel Rules Insertion	[ConnectX-5 & ConnectX-6 Dx/ BlueField & BlueField-2] Added support for a locking mechanism to enable parallel insertion of rules into the software steering using the mlx5dv_dr API. The parallel insertion improves the insertion rate and takes place when adding Rx and Tx rules via the FDB domain.	
mlx5dv_dr API Matching on Geneve Tunnel	[ConnectX-5 & ConnectX-6 Dx/ BlueField & BlueField-2] Added support for the option to match mlx5dv_dr API on Geneve tunnel using a dynamic flex parser. The option header consists of class, type, length and data. The parser should be configured using devx command, after which a rule can be created to match on parser ID and data.	
OVS-DPDK Geneve Encap/Dec ap	[ConnectX-5 & ConnectX-6 Dx/ BlueField & BlueField-2] Added support for Geneve tunneling offload, including matching on extension header.	

Category	Description	
OVS-DPDK Parallel Offloads	[ConnectX-5 & ConnectX-6 Dx/ BlueField & BlueField-2] Added support for parallel insertion and deletion of offloaded rules using multiple OVS threads.	
GTP-U TEID Modificatio n	[BlueField-2 & ConnectX-6 Dx] [Beta] Added support to modify GTP-U TEID. This support requires flex parser configuration.	
OVS-DPDK E2E Cache Support	[BlueField-2 & ConnectX-6 Dx] [Beta] Improved performance of OVS Connection Tracking flows by enabling the merge of the multi-table flow matches and actions into one joint flow.	
Tx Port Time- Stamping	[ConnectX-6 Dx and above] Transmitted packet timestamping accuracy can be improved when using a timestamp generated at the port level instead of a timestamp generated upon CQE creation. Tx port time- stamping better reflects the actual time of a packet's transmission. This feature is disabled by default. The feature can be enabled or disabled using the following command. ethtoolset-priv-flags <ifs-name> tx_port_ts on / off For further information on this feature, please see <u>Tx Port Time-Stamping</u>.</ifs-name>	
Tunnel	[ConnectX-6 Dx and above] Added support for offloading tunnel rules when the source interface is VF (in addition to uplink) in the Hypervisor.	
Rules Offload	[ConnectX-6 Dx and above] Added support for offloading tunnel rules when the source interface is OpenvSwitch bridge (internal port).	
Connection Tracking Mirroring Offload	[ConnectX-6 Dx and above] Added support for using Mirroring Offload wit Connection Tracking.	
mlx5dv_dr API ASO Flow Meter	[ConnectX-6 Dx and above] Added support for ASO flow meter using the mlx5dv_dr API, which allows for monitoring the packet rate for specific flows. When a packet hits a flow that is connected to a flow meter, the rate of packets through this meter is evaluated, and the packet is marked with a color copied into one of the C registers, according to the current rate compared to the reference rate.	
mlx5dv_dr API ASO First Hit	[ConnectX-6 Dx and above] Added support for ASO first hit using the mlx5dv_dr API, which allows for tracking rule hits by packets. When a packet hits a rule with the ASO first hit action, a flag is set indicating this event, and the original value of the flag is copied to one of the C registers.	
mlx5dv_dr API GTP-U	[ConnectX-6 Dx and above] Added mlx5dv_dr API support for matching on a new field "gtpu_first_ext_dw_0". This field enables packet filtering based	

Category	y Description		
Extension Header	on the GTP-U first extension header (first dword only). To enable parsing of tunnel GTP-U extension header, run the following command. ./cloud_fw_reset.py FLEX_PARSER_PROFILE_ENABLE=3		
IPsec Offload	[ConnectX-6 Lx and above] Added IPsec full offload support for extended sequence number, replay protection window and lifetime packet limit.		
Firmware Upgrade	[All HCAs] Firmware upgrade during MLNX_OFED installation is now done on all supported devices simultaneously rather than consecutively.		
RDMA-CM Disassociat e Support	[All HCAs] Added support for connecting kernel and RDMA-CM in a reliable way based on device index.		
New Query GID API	[All HCAs] Added support for a new query GID API that allows for querying a single GID entry by its port and GID index, or querying for all GID tables of a specific device. This API works over ioctl instead of sysfs, which accelerates the querying process.		
Multi-Host Firmware Reset	[All HCAs] Added support for performing multi-host firmware reset in order to upgrade the device firmware. Firmware reset loads the new firmware in case it was burnt on the flash and was pending activation, and reloads the current firmware image from the flash in case no new firmware was pending.		
Firmware Live Patching	[All HCAs] [Alpha] Added support for firmware live patching in the driver. Live patching updates the firmware without the need to perform firmware reset. However, it can only be applied in scenarios where the difference between the current and new firmware versions are minor, which is decided upon by the firmware itself.		
Devlink Firmware Reset	 [All HCAs] Added support in the devlink tool for performing firmware reset in order to upgrade the device firmware. Firmware reset loads the new firmware in case it was burnt on the flash and was pending activation, and reloads the current firmware image from the flash in case no new firmware was pending. For further information, please refer to the the devlink man page. Note: In order for the firmware reset to run successfully, the following conditions should be met. Each function should have the driver up and active with a version that 		
	 Each function should have the driver up and active with a version that supports this feature None of the functions has the devlink parameter enable_remote_dev_reset set to False. 		

Category	Description	
	(i) Note The current MLNX_OFED does not include the latest iproute2 version that provides support for this feature. Therefore, to be able to work with it, make sure to install the latest iproute2 version available on Github.	
Command Interface Resiliency	[All HCAs] Added a resiliency mechanism for the driver to manually poll the command event queue (EQ) in case of a command timeout. In case the resiliency mechanism finds unhandled event queue entry (EQE) due to a lost interrupt, the driver will handle it, after which the command interface returns to a healthy state.	
Offloaded Traffic Sniffer	[All HCAs] Setting a sniffer private flag is deprecated and no longer required. In order to capture offloaded/RoCE traffic, tcpdump can now be run on the RDMA device.	
Devlink Port Health Reporters	[All HCAs] Added per-port reporters to devlink health to manage per-port health activities. Users can now access the devlink port reporters by specifying the port index in addition to the device devlink name through the devlink health commands API. This update was first introduced in iproute2 v5.8. As part of this feature, mlx5e Tx and Rx reporters are now redefined as devlink port reporters. For examples, please see devlink-health manpage.	
Memory Registratio n Optimizatio n	[All HCAs] Optimized memory consumption of memory registration in huge page systems. As an example, in a 2MB huge page system, 600 MB would be saved for 100 GB memory registration.	
mlx5dv API	[All HCAs] Added support for mlx5dv API to modify the configured UDP source port for RoCE packets of a given RC/UC QP when QP is in RTS state.	
Enhanced Tx Multi- packet WQE (MPWQE)	[All HCAs] Added support for accelerating Tx datapath by saving PCI bandwidth and CPU utilization. The savings are achieved by aggregating multiple packets into a single WQE. The feature is driven by xmit_more for certain traffic types, such as UDP.	
Innova IPsec NIC Support	Removed support for the network adapter Innova IPsec (EN).	

Category	Description	
Bug Fixes	See <u>Bug Fixes</u> .	
Category	Description	
5.1-0.6.6.0		
IP-in-IP RSS Offload	[ConnectX-4 and above] Added support for receive side scaling (RSS) offload in IP-in-IP (IPv4 and IPv6).	
Devlink Port Support in Non- representor Mode	[ConnectX-4 and above] Added support for viewing the mlx5e physical devlink ports using the 'devlink port' command. This also may affect network interface names, if predictable naming scheme is configured. Suffix indicating a port number will be added to interface name.	
Devlink Health State Notification s	[ConnectX-4 and above] Added support for receiving notifications on devlink health state changes when an error is reported or recovered by one of the reporters. These notifications can be seen using the userspace 'devlink monitor' command.	
Legacy SR- IOV VF LAG Load Balancing	[ConnectX-4 and above] When VF LAG is in use, round-robin the Tx affinity of channels among the different ports, if supported by the firmware, enables all SQs of a channel to share the same port affinity. This allows the distribution of traffic sent from a VF between two ports, as well as round-robin the starting port among VFs to distribute traffic originating from single-core VMs.	
RDMA-CM DevX Support	[ConnectX-4 and above] Added support for DevX in RDMA-CM applications.	
RoCEv2 Flow Label and UDP Source Port Definition	[ConnectX-4 and above] This feature provides flow label and UDP source port definition in RoCE v2. Those fields are used to create entropy for network routes (ECMP), load balancers and 802.3ad link aggregation switching that are not aware of RoCE headers.	
RDMA Tx Steering	[ConnectX-4 and above] Enabled RDMA Tx steering flow table. Rules in this flow table will allow for steering transmitted RDMA traffic.	
Custom Parent- Domain Allocators for CQ	[ConnectX-4 and above] Enabled specific custom allocations for CQs.	

Category	Description	
mlx5dv Helper APls for Tx Affinity Port Selection	 [ConnectX-4 and above] Added support for the following mlx5dv helper APIs which enable the user application to query or set a RAW QP's Tx affinity port number in a LAG configuration. mlx5dv_query_qp_lag_port mlx5dv_modify_qp_lag_port 	
RDMA-CM Path Alignment	[ConnectX-4 and above] Added support for RoCE network path alignment between RDMA-CM message and QP data. The drivers and network components in RoCE calculate the same hash results for egress port selection both on the NICs and the switches.	
IPoIB QP Number Creation	[ConnectX-4 and above] Enabled setting the QP number of an IPoIB PKey interface in Enhanced mode. This is done using the standard ip link add command while padding the hardware address of the newly created interface. The QP number is the 2nd-4th bytes. To enable the feature, the MKEY_BY_NAME configuration should firstly be enabled in the NvConfig.	
CQ and QP Context Exposure	[ConnectX-4 and above] Exposed QP, CQ and MR context in raw format via RDMA tool.	
In-Driver xmit_more	[ConnectX-4 and above] Enabled xmit_more feature by default in kernels that lack Rx bulking support (v4.19 and above) to ensure optimized IP forwarding performance when stress from Rx to Tx flow is insufficient. In kernels with Rx bulking support, xmit_more is disabled in the driver by default, but can be enabled to achieve enhanced IP forwarding performance.	
Relaxed Ordering	 [ConnectX-4 and above] Relaxed ordering is a PCIe feature which allows flexibility in the transaction order over the PCIe. This reduces the number of retransmissions on the lane, and increases performance up to 4 times. By default, mlx5e buffers are created with Relaxed Ordering support when firmware capabilities are on and the PCI subsystem reports that CPU is not on the kernel's blocklist. Note: Some CPUs which are not listed in the kernel's blocklist may suffer from buggy implementation of relaxed ordering, in which case the user may experience a degradation in performance and even unexpected behavior. To turn off relaxed ordering and restore previous behavior, run setpci command as instructed here. Example: 	
ODP Huge Pages Support	[ConnectX-4 and above] Enabled ODP Memory Region (MR) to work with huge pages by exposing IBV_ACCESS_HUGETLB access flag to indicate that the MR range is mapped by huge pages.	

Category	Description	
	The flag is applicable only in conjunction with IBV_ACCESS_ON_DEMAND.	
Offloaded Traffic Sniffer	[ConnectX-4 and above] Removed support for Offloaded Traffic Sniffer feature and replaced its function with Upstream solution tcpdump tool.	
Connection Tracking Offload	[ConnectX-5 and above] Added support for offloading TC filters containing connection tracking matches and actions.	
Dual-Port RoCE Support	[ConnectX-5 and above] Enabled simultaneous operation of dual-port RoCE and Ethernet in SwitchDev mode.	
IP-in-IP Tunnel Offload for Checksum and TSO	[ConnectX-5 and above] Added support for the driver to offload checksum and TSO in IP-in-IP tunnels.	
Packet Pacing DevX Support	[ConnectX-5 and above] Enabled RiverMax to work over DevX with packet pacing functionality by exposing a few DV APIs from rdma-core to enable allocating/destroying a packet pacing index. For further details on usage, see man page for: mlx5dv_pp_alloc() and mlx5dv_pp_free().	
Software Steering Support for Memory Reclaiming	[ConnectX-5 and above] Added support for reclaiming device memory to the system when it is not in use. This feature is disabled by default and can be enabled using the command mlx5dv_dr_domain_set_reclaim_device_memory().	
SR-IOV Live Migration	[ConnectX-5 and above] [Beta] Added support for performing a live migration for a VM with an SR-IOV NIC VF attached to it and with minimal to no traffic disruption. This feature is supported in SwitchDev mode; enabling users to fully leverage VF TC/OVS offloads, where the failover inbox driver is in the Guest VM, and the bonding driver is in the Hyperviso Note that you must use the latest QEMU and libvirt from the Upstream <u>github.com</u> sources.	
Uplink Representor Modes	[ConnectX-5 and above] Removed support for new_netdev mode in SwitchDev mode. The new default behaviour is to always keep the NIC netdev.	
OVS-DPDK Offload Statistics	[ConnectX-5 and above] Added support for dumping connection tracking offloaded statistics.	

Category	Description	
OVS-DPDK Connection Tracking Labels Exact Matching	[ConnectX-5 and above] Added support for labels exact matching in OVS- DPDK CT openflow rules.	
OVS-DPDK LAG Support	[ConnectX-5 & ConnectX-6 Dx] Added support for LAG (modes 1,2,4) with OVS-DPDK.	
Get FEC Status on PAM4/50G	[ConnectX-6 and above] Allowed configuration of Reed Solomon and Low Latency Reed Solomon over PAM4 link modes.	
RDMA-CM Enhanced Connection Establishme nt (ECE)	[ConnectX-6 and above] Added support for allowing automatic enabling/disabling of vendor specific features during connection establishment between network nodes, which is performed over RDMA- CM messaging interface.	
RoCE Selective Repeat	[ConnectX-6 and above] This feature introduces a new QP retransmission mode in RoCE in which dropped packet recovery is done by re-sending the packet instead of re-sending the PSN window only (Go-Back-N protocol). This feature is enabled by default when RDMA-CM is being used and both connection nodes support it.	
IPsec Full Offload	[ConnectX-6 Dx & BlueField-2] [Beta] Added support for IPsec full offload (VxLAN over ESP transport).	
Hardware vDPA on OVS-DPDK	[ConnectX-6 Dx & BlueField-2] Added support for configuring hardware vDPA on OVS-DPDK. This support includes the option to fall back to Software vDPA in case the NIC installed on the driver does not support hardware vDPA.	
IPsec Crypto Offloads	[ConnectX-6 Dx] Support for IPsec Crypto Offloads feature over ConnectX-6 Dx devices and up is now at GA level.	
TLS Tx Hardware Offload	[ConnectX-6 Dx] Support for TLS Tx Hardware Offload feature over ConnectX-6 Dx devices and up is now at GA level.	
TLS Rx Hardware Offload	[ConnectX-6 Dx] [Alpha] Added support for hardware offload decryption of TLS Rx traffic over crypto-enabled ConnectX-6 Dx NICs and above.	

Category	Description	
Userspace Software Steering ConnectX-6 Dx Support	[ConnectX-6 Dx] Support for software steering on ConnectX-6 Dx adapter cards in the user-space RDMA-Core library through the mlx5dv_dr API is now at GA level.	
Kernel Software Steering ConnectX-6 Dx Support	[ConnectX-6 Dx] [Beta] Added support for kernel software steering on ConnectX-6 Dx adapter cards.	
Adapters	[ConnectX-6 Lx] Added su	pport for ConnectX-6 Lx adapter cards.
RDMA-Core Migration	[All HCAs] As of MLNX_OFED v5.1, Legacy verbs libraries have been fully replaced by RDMA-Core library. For the list of new APIs used for various MLNX_OFED features, please refer to the Migration to <u>RDMA-Core document</u> .	
Firmware Reactivation	[All HCAs] Added support for safely inserting consecutive firmware images without the need to reset the NIC in between.	
	[All HCAs] UCX-CUDA is now supported on the following OSs and platforms.	
	OS	Platform
	RedHat 7.6 ALT	PPC64LE
UCX-CUDA	RedHat 7.7	x86_64
Support	RedHat 7.8	PPC64LE/x86_64
	RedHat 7.9	x86_64
	RedHat 8.1	x86_64
	RedHat 8.2	x86_64
HCOLL- CUDA	 [All HCAs] The hcoll package includes a CUDA plugin (<u>hmca_gpu_cuda.so</u>). As of MLNX_OFED v5.1, it is built on various platforms as the package hcoll-cuda. It will be installed by default if the system has CUDA 10-2 installed. Notes: If you install MLNX_OFED from a package repository, you will need to 	

Category	Description
	 HCOLL-CUDA is supported on the same OSs that include support for UCX-CUDA (listed in the table above), except for RedHat 8.1 and 8.2.
GPUDirect Storage (GDS)	 [All HCAs] [Beta] Added support for the new technology of GDS (GPUDirect Storage) which enables a direct data path between local or remote storage, such as NFS, NVMe or NVMe over Fabric (NVMe-oF), and GPU memory. Both GPUDirect RDMA and GPUDirect Storage avoid extra copies through a bounce buffer in the CPU's memory. They enable the direct memory access (DMA) engine near the NIC or storage to move data on a direct path into or out of GPU memory, without burdening the CPU or GPU. To enable the feature, run ./mlnxofedinstallwith-nfsrdmawith-nvmfenable-gdsadd-kernel-support To get access to GDS Beta, please reach out to the GDS team at <u>GPUDirectStorageExt@nvidia.com</u>. For the list of operating systems on which GDS is supported, see here.

Category	Description	
5.0-2.1.8.0		
Kernel Software Managed Flow Steering (SMFS) Performanc e	[ConnectX-5 and above] Improved the performance of Kernel software steering by reducing its memory consumption.	
NEO-Host SDK	[All HCAs] Added support for NEO-Host SDK installation on MLNX_OFED.	
Bug Fixes	See <u>Bug Fixes</u> .	
5.0-1.0.0.0	5.0-1.0.0.0	
Adapters	[ConnectX-6 Dx] Added support for ConnectX-6 Dx adapter cards.	
Userspace Software Steering ConnectX-6 Dx Support	[ConnectX-6 Dx] [Beta] Added support for software steering on ConnectX-6 Dx adapter cards in the user-space RDMA-Core library through the mlx5dv_dr API.	

Category	Description
Virtual Output Queuing (VoQ) Counters	[ConnectX-6 Dx and above] Exposed rx_prio[p]_buf_discard, rx_prio[p]_wred_discard and rx_prio[p]_marked firmware counters that count the number of packets that were dropped due to insufficient resources.
IPsec Crypto Offloads	[ConnectX-6 Dx and above] [Beta] IPsec crypto offloads are now supported on ConnectX-6 Dx devices and up. The offload functions use the existing ip xfrm tool to activate offloads on the device. It supports transport/tunnel mode with AES-GCM IPsec scheme.
TLS TX Hardware Offload	[ConnectX-6 Dx and above] [Alpha] Added support for hardware offload encryption of TLS traffic. Note: Not supported in ConnectX-6 Lx adapter cards.
VirtIO Acceleration through Datapath I/O Processor (vDPA)	 [ConnectX-6 Dx and above] Added support to enable mapping the VirtIO access region (VAR) to be used for doorbells by vDPA applications. Specifically, the following DV APIs were introduced (see man page for more details): mlx5dv_alloc_var() mlx5dv_free_var()
	Note: Not supported in ConnectX-6 Lx adapter cards.
Resource Allocation on External Memory	[ConnectX-5 and above] Added support to enable overriding mlx5 internal allocations in order to let applications allocate some resources on external memory, such as that of the GPU. The above is achieved by extending the parent domain object with custom allocation callbacks. Currently supported verbs objects are: QP, DBR, RWQ, SRQ.
Hardware Clock Exposure	[ConnectX-5 and above] Added support for querying the adapter clock via mlx5dv_query_device.
ODP Diagnostic Counters	 [ConnectX-5 and above] Added ODP diagnostics counters for the following items per MR (memory region) within IB/mlx5 driver: 1. Page faults: Total number of faulted pages. 2. Page invalidations: Total number of pages invalidated by the OS during all invalidation events. The translations can no longer be valid due to either non-present pages or mapping changes. 3. Prefetched pages: When prefetching a page, a page fault is generated in order to bring the page to the main memory.

Category	Description	
Devlink Health CR- Space Dump	[ConnectX-5 and above] Added the option to dump configuration space via the devlink tool in order to improve debug capabilities.	
Multi-packet TX WQE Support for XDP Transmit Flows	[ConnectX-5 and above] The conventional TX descriptor (WQE or Work Queue Element) describes a single packet for transmission. Added driver support for the HW feature of multi-packet TX WQEs in XDP transmit flows. With this, the HW becomes capable of working with a new and improved WQE layout that describes several packets. In effect, this feature saves PCI bandwidth and transactions, and improves transmit packet rate.	
OVS-Kernel ToS Rewrite	[ConnectX-5 and above] Added support for Type of Service (ToS) rewrite in the OVS-Kernel.	
OVS-Kernel Mirroring	[ConnectX-5 and above] Added support for mirroring output in SwitchDev mode in the OVS-Kernel. The mirroring port may either be a local or a remote VF, using VxLAN or GRE encapsulations.	
GENEVE Encap/Deca p Rules Offload	[ConnectX-5 and above] Added support for GENEVE encapsulation/decapsulation rules offload.	
GPRS Tunneling Protocol (GTP) Header	[ConnectX-5 and above] [Beta] Added support for matching (filtering) GTP header-based packets using mlx5dv_dr API over user-space RDMA- Core library.	
Multi Packet Tx WQE Support for XDP Transmit Flows	[ConnectX-5 and above] Added driver support for the hardware feature of multi-packet Tx to work with a new and improved WQE layout that describes several packets instead of a single packet for XDP transmission flows. This saves PCI bandwidth and transactions, and improves transmit packet rate.	
Userspace Software Steering Debugging API	[ConnectX-5 and above] [Beta] Added support for software steering to dump flows for debugging purposes in the user-space RDMA-Core library through the mlx5dv_dr API.	
Kernel Software Steering for	[ConnectX-5 and above] [Beta] Added support for updating CT rules using the software steering mechanism.	

Category	Description
Connection Tracking (CT)	
Kernel Software Steering Remote Mirroring	[ConnectX-5 and above] [Beta] Added support for updating remote mirroring rules using the software steering mechanism.
OVS-DPDK Support	[ConnectX-5 and BlueField] Added OVS-DPDK component as part of the MLNX_OFED package with hardware offload capabilities.
OVS-DPDK Connection Tracking	[ConnectX-5 and BlueField] [Beta] Added support for OvS-DPDK Connection Tracking hardware offload.
OVS-DPDK VirtIO Acceleration through VF Relay	[ConnectX-5 and BlueField] Added support for OVS-DPDK VirtlO Acceleration through VF Relay (also known as Software vDPA) forwarding of traffic from VF to Virtio VM and vice-versa.
OVS-DPDK VXLAN Encap/Deca p	[ConnectX-5 and BlueField] Added support for OVS-DPDK VXLAN encapsulation and decapsulation hardware offload.
Discard Counters	[ConnectX-4 and above] Exposed rx_prio[p]_discards discard counters per priority that count the number of received packets dropped due to lack of buffers on the physical port.
MPLS Traffic	[ConnectX-4 and above] Added support for reporting TSO and CSUM offload capabilities for MPLS tagged traffic and, allowed the kernel stack to use these offloads.
mlx5e Max Combined Channels	[ConnectX-4 and above] Increased the driver's maximal combined channels value from 64 to 128 (however, note that OOB value will not cross 64). 128 is the upper bound. Lower maximal value can be seen on the host, depending on the number of cores and MSIX's configured by the firmware.
RoCE Accelerator Counters	 [ConnectX-4 and above] Added the following RoCE accelerator counters: roce_adp_retrans - counts the number of adaptive retransmissions for RoCE traffic roce_adp_retrans_to - counts the number of times RoCE traffic reached timeout due to adaptive retransmission

Category	Description
	 roce_slow_restart - counts the number of times RoCE slow restart was used roce_slow_restart_cnps - counts the number of times RoCE slow restart generated CNP packets roce_slow_restart_trans - counts the number of times RoCE slow restart changed state to slow restart
Migration to RDMA-Core	 [All HCAs] The default installation of the userspace is now the RDMA-Core library instead of the legacy verbs. This achieves most of the legacy experimental verbs' functionalities, and more. For NVIDIA VMA or NVIDIA RiverMax, use experimental verbs (prefix "ibv_exp"). For further information on the migration to RDMA-Core and the list of new APIs used for various MLNX_OFED features, please refer to the Migration to RDMA-Core document.
ibdev2netde v Tool Output	[All HCAs] ibdev2netdev tool output was changed such that the bonding device now points at the bond instead of the slave interface.
Memory Region	[All HCAs] Added support for the user to register memory regions with a relaxed ordering access flag. This can enhance performance, depending on architecture and scenario.
Devlink Health Reporters	[All HCAs] Added support for monitoring and recovering from errors that occur on the RX queue, such as cookie errors and timeout.
GSO Optimization	[All HCAs] Improved GSO (Generic Segmentation Offload) workload performance by decreasing doorbells usage to the minimum required.
TX CQE Compressio n	[All HCAs] Added support for TX CQE (Completion Queue Element) compression. Saves on outgoing PCIe bandwidth by compressing CQEs together. Disabled by default. Configurable via private flags of ethtool.
Firmware Versions Query via Devlink	[All HCAs] Added the option to query for running and stored firmware versions using the devlink tool.
Firmware Flash Update via Devlink	[All HCAs] Added the option to update the firmware image in the flash using the devlink tool. Usage: devlink dev flash <dev> file <file_name>.mfa2 For further information on how to perform this update, see "Updating Firmware Using ethtool/devlink and .mfa2 File" section in MFT User Manual.</file_name></dev>

Category	Description
Devlink Health WQE Dump	[All HCAs] Added support for WQE (Work Queue Element) dump, triggered by an error on Rx/Tx reporters. In addition, some dumps (not triggered by an error) can be retrieved by the user via devlink health reporters.
GENEVE Tunnel Stateless Offload	[All HCAs] Added support for GENEVE tunneled hardware offloads of TSO, CSUM and RSS.
TCP Segmentatio n and Checksum Offload	[All HCAs] Added TCP segmentation and checksum offload support for MPLS-tagged traffic.

Category	Description	
4.7-3.2.9.0	4.7-3.2.9.0	
Uplink Representor Modes	[ConnectX-5 and above] Added support for new_netdev and nic_netdev uplink representor modes. For further information on how to configure these modes, please refer to Configuring Uplink Representor Mode.	
	[ConnectX-5 and above] Added new mlx5_core module parameter "num_of_groups", which controls the number of large groups in the FDB flow table. Note : The default value of num_of_groups may change per MLNX_OFED driver version. The following table lists the values that must be set when upgrading the MLNX_OFED version prior to driver load , in order to achieve the same OOB experience.	
mlx5_core	MLNX_OFED Version	num_of_groups Default Value
	v4.7-3.2.9.0	4
	v4.6-3.1.9.0.14	15
	v4.6-3.1.9.0.15	15
	v4.5-1.0.1.0.19	63
	For further information, please Traffic Patterns section in MLN	e refer to Performance Tuning Based on NX_OFED User Manual.

Category	Description	
VFs Groups Minimum Bandwidth Rate	[ConnectX-5] Added support for setting a minimum bandwidth rate on a group of VFs (BW guarantee) to ensure this group is able to transmit at least the amount of bandwidth specified on the wire.	
Direct Verbs Support for Batch Counters on Root Table	[ConnectX-5] Added support for mlx5dv_dr API to set batch counters for root tables.	
Modify Header	[ConnectX-5 and BlueField] Added support for mlx5dv_dr_actions to support up to 32 modify actions.	
mlx5dv_dr Memory Consumptio n	[ConnectX-5 and BlueField] Reduced the mlx5dv_dr API memory consumption by improving the memory allocator.	
mlx5dv_dr Memory Allocation	[ConnectX-5 and BlueField] Reduced memory allocation time when using the mlx5dv_dr API. This is particularly significant for the first inserted rules on which memory is allocated.	
Mediated Devices	[ConnectX-5 and BlueField] Added support for mediated devices that allows the creation of accelerated devices without SR-IOV on the Bluefield® system. For further information on mediated devices and how to configure them, please refer to Mediated Devices section in MLNX_EN User Manual.	
4.7-1.0.0.1	·	
Counters Monitoring	[ConnectX-4 and above] Added support for monitoring selected counters and generating a notification event (Monitor_Counter_Change event) upon changes made to these counters. The counters to be monitored are selected using the SET_MONITOR_COUNTER command.	
Signature Offload Kernel Verbs Enhanceme nts	[ConnectX-4 and above] Added a new API which enables posting a single WR that completes the Protection Information (PI) operation internally. This reduces CPU utilization for posting and processing multiple WRs and improves performance by choosing the optimal mkey for the hardware according to the buffer memory layout.	
EEPROM Device Thresholds via Ethtool	[ConnectX-4 and above] Added support to read additional EEPROM information from high pages of modules such as SFF-8436 and SFF-8636 Such information can be: 1. Application Select table 2. User writable EEPROM 3. Thresholds and alarms - Ethtool dump works on active cables	

Category	Description
	only (e.g. optic), but thresholds and alarms can be read with "offset" and "length" parameters in any cable by running: ethtool -m <devname> offset X length Y</devname>
Performanc e Improvemen ts	 [ConnectX-4 and above] Updated Blueflame capability reporting to prevent redundant use of Blueflame when Write-combining is not supported. Added Blueflame capabilities over VFs.
RDMA_RX RoCE Steering Support	[ConnectX-4 and above] Added the ability to create rules to steer RDMA traffic, with two destinations supported: DevX object and QP. Multiple priorities are also supported.
SRQ and XRC Support on On Demand Paging (ODP) Memory Region (MR)	[ConnectX-4 and above] Added support for using ODP MR with SRQ WQEs and XRC transport.
Indirect Mkey ODP	[ConnectX-4 and above] Added the ability to create indirect Mkeys with ODP support over DevX interface.
DevX Asynchrono us Query Commands	[ConnectX-4 and above] Added support for running QUERY commands over the DevX interface in an asynchronous mode. This enables applications to issue many commands in parallel while firmware processes the commands.
Implicit ODP	[ConnectX-4 and above] Added support for reporting implicit ODP support to user applications in order to allow better granularity over ODP creation.
Devlink Health Utility	[ConnectX-4 and above] Added support for real-time alerting of functionality issues that may be found in a system component (reporter). This utility helps detect and recover from a problem with a PCI device. It provides a centralize status of drivers' health activities in the generic Devlink instance and inter alia, supports the following:
	 Storing real-time error dumps Performing automatic (configurable) real-time reporter recovery Performing real-time reporter diagnosis Indicating real-time reporter's health status

Category	Description
	 Providing admins with the ability to dump, diagnose and recover a reporter Providing admins with the ability to configure a reporter
User-Mode Memory Registration (UMR)	[ConnectX-4 and above] Enabled registration of memory patterns that can be used for future RDMA operations.
GENEVE Tunnel Stateless Offload	[ConnectX-4 and above] Added support for Generic Network Virtualization Encapsulation (GENEVE) tunneled hardware offload of TSO, CSUM and RSS.
ODP Pre- fetch	[ConnectX-4 and above] Added support for pre-fetching a range of an on- demand paging (ODP) memory region (MR), this way reducing latency by making pages present with RO/RW permissions before the actual IO is conducted.
Fragmented QPs Buffer	[ConnectX-4 and above] Added the ability to allocate a fragmented buffer to in-kernel QP creation requests, in cases of large QP size requests that used to fail due to low memory resources on the host.
Flow Counters Batch Query	[ConnectX-4 and above] Allowed flow counters created with the DevX interface to be attached to flows created with the raw flow creation API.
DevX Privilege Enforcemen t	[ConnectX-4 and above] Enforced DevX privilege by firmware. This enables future device functionality without the need to make driver changes unless a new privilege type is introduced.
DevX Interoperabil ity APIs	[ConnectX-4 and above] Added support for modifying and/or querying for a verb object (including CQ, QP, SRQ, WQ, and IND_TBL APIs) via the DevX interface. This enables interoperability between verbs and DevX.
Counters Monitoring	[ConnectX-4 and above] Added support for monitoring selected counters and generating a notification event (Monitor_Counter_Change event) upon changes made to these counters. The counters to be monitored are selected using the SET_MONITOR_COUNTER command.

Category	Description
Rx Hash Fields Configuratio n	[ConnectX-4 and above] Added the ability to configure Rx hash fields used for traffic spreading into Rx queues using ETHTOOL_SRXFH and ETHTOOL_GRXFH ethtool commands. Built-in Receive Side Scaling (RSS) profiles can now be changed on the following traffic types: UDP4, UDP6, TCP4 and TCP6. This configuration affects both outer and inner headers.
Equal Cost Multi-Path (ECMP)	[ConnectX-4 Lx and above] Added support for offloading ECMP rules by tracking software multipath route and related next-hops, and reflecting this as port affinity to the hardware.
VF LAG	[ConnectX-4 Lx and above] Added support for High Availability and load balancing for Virtual Functions of different physical ports in SwitchDev SR-IOV mode.
Uplink Representor s	[ConnectX-4 Lx and above] Exposed PF (uplink) representors in SwitchDev mode, similarly to VF representors, as an infrastructure improvement for SmartNICs.
Userspace Software Steering for eSwitch	[ConnectX-5] Added software steering capabilities to the SR-IOV eSwitch. Software steering enables better rules insertion rate compared to the current firmware-based solution. This is achieved by performing calculations on the main CPU which allows for higher insertion rates.
Userspace Software Steering for NICs	[ConnectX-5] Added software steering capabilities to NIC Rx/Tx. Software steering enables better rules insertion rate compared to the current firmware-based solution. This is achieved by performing calculations on the main CPU which allows for higher insertion rates. This solution was designed to work with Virtio DPDK. Note : Support will be enabled by default once the support for GID change is added.
ASAP ²	[ConnectX-5 and above] Incorporated the documentation of <i>Accelerated Switching And Packet Processing (ASAP²): Hardware Offloading for vSwitches</i> into MLNX_OFED Release Notes and User Manual.
QP Counters and Firmware Errors per PID	[ConnectX-5 and above] QP counters and flow counters are now set per Process ID (PID) to allow better visibility of RDMA error states. Users will be able to manually tune the Q counter to monitor specific QPs, or automatically monitor QPs according to predefined criteria, such as the QP type.
ODP over DC	[ConnectX-5 and above] Added support for On-Demand Paging (ODP) over DC transport.
Address Translation Services	[ConnectX-5 and above] Added support for Address Translation Services (ATS) feature, which improves performance for virtualized PeerDirect

Category	Description
	applications by caching PA-> MA translations and preventing PCI transactions from going to the root complex.
XDP Inline Transmissio n of Small Packets	[ConnectX-5 and above] Added support for when forwarding packets with XDP, a packet smaller than 256 bytes would be sent inline within its WQE Tx descriptor for better performance. The number of packets that are transmitted inline depends on CPUs load, where lower load leads to a higher number of inline transmission.
VLAN Rewrite	[ConnectX-5 and above] Added support for offloading VLAN ID modify operation, allowing the user to replace the VLAN tag of the incoming frame with a user-specified VLAN tag value.
CQE Padding	[ConnectX-5 and above] Added support for padding 64B CQEs to 128B cache lines to improve performance on 128B cache line systems, such as PPC.
XDP Multi- Packet Tx Work Queue Element (WQE)	[ConnectX-5 and above] Added support for Multi-Packet Tx WQEs in XDP transmit flows to work with a new and improved WQE layout that describes several packets.This saves PCI bandwidth and transactions, and improves transmit packet rate.
ConnectX Device IDs	 [ConnectX-6] Added support for the following new device IDs: ConnectX-6 Dx (PF) ConnectX Family mlx5Gen Virtual Function (VF) Note that every new device (adapter) VF will be identified with this device ID. Different VF models will be distinguished by their revision ID.
Ethtool 200Gbps	[ConnectX-6 and above] ConnectX-6 hardware introduces support for 200Gbps and 50Gbps-per-lane link mode. The driver supports full backward compatibility with previous configurations. Note that in order to advertise newly added link-modes, the full bitmap related to the link modes must be advertised from ethtool man page. NOTE : This feature is firmware-dependent. Currently, ConnectX-6 Ethernet firmware supports up to 100Gbps only. Thus, this capability may not function properly using the current driver and firmware versions.
HDR Link Speed Exposure	[ConnectX-6 and above] Added support for HDR link speed in CapabilityMask2 field in port attributes.
QP Packet Based Credit	[ConnectX-6 and above] Added support for an alternative end-to-end credit mode for QP creation. Credits transported from the responder to

Category [Description
5	the requester are now issued per packet. This is particularly useful for sending large RDMA messages from HCA to switches that are short in memory.
Device Emulation Infrastructur e	[BlueField] Added support for Device Emulation in BlueField. This mechanism allows function-A to perform operations on behalf of function-B. The emulation manager creates a channel (named VHCA_TUNNEL general object) that acts as the direct command interface between the emulated function host and the HCA hardware. The emulation software creates this tunnel for every managed function and issues commands via the DevX general command interface.
Migration to	[All HCAs] Legacy verbs remain the default userspace installation option in the MLNX_OFED. However, as of MLNX_OFED v4.7, you can opt to install full RDMA-Core based userspace by adding the upstream-libs flag to the mInxofedinstall script.
Installation F via I	[All HCAs] The repository providing legacy verbs has been moved from RPMS or DEBS folders to RPMS/MLNX_LIBS and DEBS/MLNX_LIBS. In addition, a new repository providing RDMA-Core based userspace has been added to RPMS/UPSTREAM_LIBS and DEBS/UPSTREAM_LIBS.
	[All HCAs] Added support for NFS over RDMA (NFSoRDMA) module over the OSs listed in <u>NFSoRDMA Supported OSs</u> section.
	[All HCAs] As of MLNX_OFED v4.7, NFSoRDMA driver is no longer installed by default. In order to install it over a supported kernel, add the "with- nfsrdma" installation option to the "mInxofedinstall" script.
	[All HCAs] Added a new option to rdma_set_option that allows applications to override the RDMA-CM's QP ACK timeout value.
· ·	[All HCAs] Added a unique ID for each verbs object to allow direct query over rdma-tool and rdma-netlink for enhanced debuggability.
I ANNIICATION I	[All HCAs] Added support for the RDMA application to manage its own QPs and use RDMA-CM only for exchanging Address information.
Bug Fixes	See <u>Bug Fixes</u> .

Category	Description
4.6-1.0.1.0	
Devlink Configuratio	[ConnectX-3/ConnectX-3 Pro] Added support for a set of configuration parameters that can be changed by the user through the Devlink user

Category	Description
n Parameters Tool	interface.
ODP Pre- fetch	[ConnectX-4 and above] Added support for pre-fetching a range of an on-demand paging (ODP) memory region (MR), this way reducing latency by making pages present with RO/RW permissions before the actual IO is conducted.
DevX Privilege Enforcement	[ConnectX-4 and above] Enforced DevX privilege by firmware. This enables future device functionality without the need to make driver changes unless a new privilege type is introduced.
DevX Interoperabil ity APIs	[ConnectX-4 and above] Added support for modifying and/or querying for a verb object (including CQ, QP, SRQ, WQ, and IND_TBL APIs) via the DevX interface. This enables interoperability between verbs and DevX.
DevX Asynchrono us Query Commands	[ConnectX-4 and above] Added support for running QUERY commands over the DevX interface in an asynchronous mode. This enables applications to issue many commands in parallel while firmware processes the commands.
DevX User- space PRM Handles Exposure	[ConnectX-4 and above] Exposed all PRM handles to user-space so DevX user application can mix verbs objects with DevX objects. For example: Take the cqn from the created ibv_cq and use it on a devx)create(QP).
Indirect Mkey ODP	[ConnectX-4 and above] Added the ability to create indirect Mkeys with ODP support over DevX interface.
XDP Redirect	[ConnectX-4 and above] Added support for XDP_REDIRECT feature for both ingress and egress sides. Using this feature, incoming packets on one interface can be redirected very quickly into the transmission queue of another capable interface. Typically used for load balancing.
RoCE Disablement	[ConnectX-4 and above] Added the option to disable RoCE traffic handling. This enables forwarding of traffic over UDP port 4791 that is handled as RoCE traffic when RoCE is enabled. When RoCE is disabled, there is no GID table, only Raw Ethernet QP type is supported and RoCE traffic is handled as regular Ethernet traffic.
Forward Error Correction (FEC) Encoding	[ConnectX-4 and above] Added the ability to query and modify Forward Error Correction (FEC) encoding, as well as disabling it via Ethtool.

Category	Description
RAW Per- Lane Counters Exposure	[ConnectX-4 and above] Exposed RAW error counters per cable-module lane via ethtool stats. The counters show the number of errors before FEC correction (if enabled). For further information, please see phy_raw_errors_lane[i] under Physical Port Counters section in <u>Understanding mlx5 ethtool Counters</u> Community post.
VF LAG	[ConnectX-4 Lx and above] Added support for High Availability and load balancing for Virtual Functions of different physical ports in SwitchDev SR-IOV mode.
ASAP ² Offloading VXLAN Decapsulatio n with HW LRO	[ConnectX-5 and above] Added support for performing hardware Large Receive Offload (HW LRO) on VFs with HW-decapsulated VXLAN. For further information on the VXLAN decapsulation feature, please refer to ASAP ² User Manual under <u>nvidia.com/en-us/networking/</u> \rightarrow Products \rightarrow Software \rightarrow ASAP ² .
PCI Atomic Operations	[ConnectX-5 and above] Added the ability to run atomic operations on local memory without involving verbs API or compromising the operation's atomicity.
Equal-Cost Multi-Path (ECMP) Routing Offloading	[ConnectX-5 and above] Enabled Equal-Cost Multi-Path (ECMP) Routing offloading. Equal-Cost Multi-Path (ECMP) is a forwarding mechanism for routing packets along multiple paths of equal cost with the goal to achieve almost equally distributed link load sharing.
VXLAN over VLAN	[ConnectX-5 and above] VXLAN over VLAN enables the user to use VXLAN offloads' benefit to offload VLAN tagged tunnels thus boost system's performance.
VLAN Rewrite	[ConnectX-5 and above] Rewriting VLAN tags allows the user to replace the VLAN tag of the incoming frame with a user-specified VLAN tag value.
Virtual Ethernet Port Aggregator (VEPA)	[ConnectX-5] Added support for activating/deactivating Virtual Ethernet Port Aggregator (VEPA) mode on a single virtual function (VF). To turn on VEPA on the second VF, run: bridge link set dev <netdev> hwmode vepa</netdev>
VFs Rate Limit	[ConnectX-5] Added support for setting a rate limit on groups of Virtual Functions rather on an individual Virtual Function.
ConnectX-6 Support	[ConnectX-6] [Beta] Added support for ConnectX-6 (VPI only) adapter cards.

Category	Description
	NOTE : In HDR installations that are built with remotely managed Quantum-based switches, the switch's firmware must be upgraded to version 27.2000.1142 prior to upgrading the HCA's (ConnectX-6) firmware to version 20.25.1500. When using ConnectX-6 HCAs with firmware v20.25.1500 and connecting them to Quantum-based switches, make sure the Quantum firmware version is 27.2000.1142 in order to avoid any critical link issues.
Ethtool 200Gbps	 [ConnectX-6] ConnectX-6 hardware introduces support for 200Gbps and 50Gbps-per-lane link mode. MLNX_OFED supports full backward compatibility with previous configurations. Note that in order to advertise newly added link-modes, the full bitmap related to the link modes must be advertised from ethtool man page. For the full bitmap list per link mode, please refer to MLNX_OFED User Manual. NOTE: This feature is firmware-dependent. Currently, ConnectX-6 Ethernet firmware supports up to 100Gbps only. Thus, this capability may not function properly using the current driver and firmware versions.
PCle Power State	 [ConnectX-6] Added support for the following PCIe power state indications to be printed to dmesg: 1. Info message #1: PCIe slot power capability was not advertised. 2. Warning message: Detected insufficient power on the PCIe slot (xxxW). 3. Info message #2: PCIe slot advertised sufficient power (xxxW). When indication #1 or #2 appear in dmesg, user should make sure to use a PCIe slot that is capable of supplying the required power.
Message Signaled Interrupts-X (MSI-X) Vectors	[mlx5] Added support for using a single MSI-X vector for all control event queues instead of one MSI-X vector per queue in a virtual function driver. This frees extra MSI-X vectors to be used for completion event queue, allowing for additional traffic channels in the network device.
Send APIs	[mlx5] Introduced a new set of QP Send operations (APIs) which allows extensibility for new Send opcodes.
DC Data- path	[mlx5] Added DC QP data-path support using new Send APIs introduced in Direct Verbs (DV).
BlueField Support	[BlueField] BlueField is now fully supported as part of the NVIDIA OFED mainstream version sharing the same code baseline with all the adapters product line.

Category	Description
Representor Name Change	 [BlueField] In SwitchDev mode: Uplink representors are now called p0/p1 Host PF representors are now called pf0hpf/pf1hpf VF representors are now called pf0vfN/pf1vfN
ECPF Net Devices	[BlueField] In SwitchDev mode, net devices enp3s0f0 and enp3s0f1 are no longer created.
Setting Host MAC and Tx Rate Limit from ECPF	[BlueField] Expanded to support VFs as well as the host PFs.
RDMA-CM Application Managed QP	[All HCAs] Added support for the RDMA application to manage its own QPs and use RDMA-CM only for exchanging Address information.
RDMA-CM QP Timeout Control	[All HCAs] Added a new option to rdma_set_option that allows applications to override the RDMA-CM's QP ACK timeout value.
MLNX_OFED Verbs API	[All HCAs] As of MLNX_OFED v5.0 release (Q1 2020) onwards, MLNX_OFED Verbs API will be migrated from the legacy version of the user space verbs libraries (libibervs, libmlx5) to the upstream version rdma-core. More details are available in MLNX_OFED user manual under Installing Upstream rdma-core Libraries.
Bug Fixes	See <u>Bug Fixes</u> .

Category	Description
4.5-1.0.1.0	
VFs per PF	[ConnectX-5] Increased the amount of maximum virtual functions (VF) that can be allocated to a physical function (PF) to 127 VF.
SW-Defined UDP Source Port for RoCE v2	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] UDP source port for RoCE v2 packets is now calculated by the driver rather than the firmware, achieving better distribution and less congestion. This mechanism works for RDMA- CM QPs only, and ensures that RDMA connection messages and data messages have the same UDP source port value.
Local Loopback	[mlx5 Driver] Added the ability to manually disable Local Loopback regardless of the number of open user-space transport domains.

Category	Description
Disable	
Adapter Cards	[ConnectX-6] Added support for ConnectX-6 Ready. For further information, please contact support <u>networking-support@nvidia.com</u> .
NEO-Host	[All HCAs] Integrated NEO-Host for orchestration and management of host networking into MLNX_OFED package.
Bug Fixes	See <u>Bug Fixes</u> .
4.4-2.0.7.0	
Bug Fixes	See <u>Bug Fixes</u> .
4.4-1.0.0.0	
Adaptive Interrupt	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for adaptive Tx, which optimizes the moderation values of the Tx CQs on runtime for maximum throughput with minimum CPU overhead. This mode is enabled by default.
Moderation	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Updated Adaptive Rx to ignore ACK packets so that queues that only handle ACK packets remain with the default moderation.
Docker Containers [Beta]	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for Docker containers to run over Virtual RoCE and InfiniBand devices using SR-IOV mode.
VF Statistics	 [ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Performed the following virtual function statistics changes: Added tx_broadcast and tx_multicast counters Included RDMA statistics for existing counters
Force TTL	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for setting a global TTL value for all RC QPs and rdma-cm QPs.
Firmware Tracer	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added a new mechanism for the device's FW/HW to log important events into the event tracing system (/sys/kernel/debug/tracing) without requiring any NVIDIA-specific tool. Note: This feature is enabled by default.
CR-Dump	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Accelerated the original cr- dump by optimizing the reading process of the device's CR-Space snapshot.

Category	Description
RoCE ICRC Error Counter	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for a new counter that exposes the amount of corrupted RoCE packets that arrive with bad Invariant Cyclic Redundancy Code (ICRC).
VST Q-in-Q	[ConnectX-4/ConnectX-4 Lx] Added support for C-tag (0x8100) VLAN insertion to tagged packets in VST mode.
Ethernet Tunneling Over IPoIB Driver (eIPoIB)	[ConnectX-4] Re-added support for eth_ipoib driver, which provides a standard Ethernet interface to be used as a Physical Interface (PIF) into the Hypervisor virtual network, and serves one or more Virtual Interfaces (VIF).
OVS Offload using ASAP ²	[ConnectX-4 Lx/ConnectX-5] Added support for NVIDIA Accelerated Switching And Packet Processing (ASAP ²) technology, which allows OVS offloading by handling OVS data-plane, while maintaining OVS control- plane unmodified. OVS Offload using ASAP ² technology provides significantly higher OVS performance without the associated CPU load. For further information, refer to ASAP ² Release Notes under nvidia.com/en-us/networking/.com \rightarrow Products \rightarrow Software \rightarrow ASAP ² .
Upstream Libraries	[All HCAs] Added a repository repodata to support installing upstream libraries (based on upstream rdma-core), using he Operating System's standard package manager (yum, apt-get, etc.). For further information, please refer to "Installing Upstream rdma-core Libraries" section in MLNX_OFED User Manual Note: This is intended only for DPDK users.
Installation	[All HCAs] Added support for new metadata packages that only install userspace packages at a time (without any kernel packages), using the Operating System's standard package manager (yum, apt-get, etc.). These metadata packages will have the suffix "-user-only". For example: "mInx- ofed-all-user-only".
Bug Fixes	See <u>Bug Fixes</u> .
4.3-1.0.1.0	
Adaptive Interrupt Moderation	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for adaptive Tx, which optimizes the moderation values of the Tx CQs on runtime for maximum throughput with minimum CPU overhead. This mode is enabled by default.
	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Updated Adaptive Rx to ignore ACK packets so that queues that only handle ACK packets remain with the default moderation.

Category	Description
Docker Containers [Beta]	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for Docker containers to run over Virtual RoCE and InfiniBand devices using SR-IOV mode.
VF Statistics	 [ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Performed the following virtual function statistics changes: Added tx_broadcast and tx_multicast counters Included RDMA statistics for existing counters
Force TTL	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for setting a global TTL value for all RC QPs and rdma-cm QPs.
Firmware Tracer	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added a new mechanism for the device's FW/HW to log important events into the event tracing system (/sys/kernel/debug/tracing) without requiring any NVIDIA-specific tool. Note: This feature is enabled by default.
CR-Dump	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Accelerated the original cr- dump by optimizing the reading process of the device's CR-Space snapshot.
RoCE ICRC Error Counter	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for a new counter that exposes the amount of corrupted RoCE packets that arrive with bad Invariant Cyclic Redundancy Code (ICRC).
VST Q-in-Q	[ConnectX-4 andConnectX-4 Lx] Added support for C-tag (0x8100) VLAN insertion to tagged packets in VST mode.
Ethernet Tunneling Over IPoIB Driver (eIPoIB)	[ConnectX-4] Re-added support for eth_ipoib driver, which provides a standard Ethernet interface to be used as a Physical Interface (PIF) into the Hypervisor virtual network, and serves one or more Virtual Interfaces (VIF).
OVS Offload using ASAP2	[ConnectX-4 Lx/ConnectX-5] Added support for NVIDIA Accelerated Switching And Packet Processing (ASAP2) technology, which allows OVS offloading by handling OVS data-plane, while maintaining OVS control- plane unmodified. OVS Offload using ASAP2 technology provides significantly higher OVS performance without the associated CPU load. For further information, refer to ASAP ² Release Notes under <u>nvidia.com/en-us/networking/.com</u> \rightarrow Products \rightarrow Software \rightarrow ASAP ² .
Upstream Libraries	[All HCAs] Added a repository repodata to support installing upstream libraries (based on upstream rdma-core), using the Operating System's

Category	Description
	standard package manager (yum, apt-get, etc.). For further information, please refer to "Installing Upstream rdma-core Libraries" section in MLNX_OFED User Manual Note: This is intended only for DPDK users.
Installation	[All HCAs] Added support for new metadata packages that only install userspace packages at a time (without any kernel packages), using the Operating System's standard package manager (yum, apt-get, etc.). These metadata packages will have the suffix "-user-only". For example: "mlnx- ofed-all-user-only".
Bug Fixes	See <u>Bug Fixes</u> .
4.3-1.0.1.0	
Multi-Packet	[ConnectX-5] Added support for the following multi-packet WR related verbs for control path:
Work Request	ibv_exp_query_deviceibv_exp_create_srq
(WR)	For further information on the use of these verbs, please refer to the Verbs man page.
Coherent Accelerator Processor Interface (CAPI) [beta]	[ConnectX-5] Added support for CAPI, an interface that enables ConnectX-5 adapter cards to provide the best performance for Power and OpenPower based platforms.
Tunneled Atomic	[ConnectX-5] Added support for RDMA atomic commands offload so that when an RDMA Write operation is issued, the payload indicates which atomic operation to perform, instead of being written to the Memory Region (MR).
Packet Pacing	 [ConnectX-5] Added support for the following advanced burst control parameters: max_burst_sz - for indicating the maximal burst size of packets typical_pkt_sz - for improving the accuracy of the rate limiter
Erasure Coding Offload verbs	[ConnectX-5] Added support for erasure coding offload software verbs (encode/decode/update API) supporting a number of redundancy blocks (m) greater than 4.

Category	Description
Virtual MAC	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Removed support for Virtual MAC feature.
RoCE LAG	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added out of box RoCE LAG support for RHEL 7.2 and RHEL 6.9.
Dropped	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added a new counter <i>rx_steer_missed_packets</i> which provides the number of packets that were received by the NIC, yet were discarded/dropped since they did not match any flow in the NIC steering flow table.
Counters	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added the ability for SR-IOV counter <i>rx_dropped</i> to count the number of packets that were dropped while vport was down.
Relaxed Ordering (RSYNC)	[ConnectX-4/ ConnectX-4 Lx/ConnectX-5] Added support for RSYNC feature to ensure correct ordering of memory operations between the GPU and HCA.
Reset Flow	[mlx5 Driver] Added support for triggering software reset for firmware/driver recovery. When fatal errors occur, firmware can be reset and driver reloaded.
Striding RQ with HW Time- Stamping	[ConnectX-4 Lx/ConnectX-5] Added the option to retrieve the HW timestamp when polling for completions from a completion queue that is attached to a multi-packet RQ (Striding RQ).
4.2-1.2.0.0	
DSCP Trust Mode	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for automatically setting the number of TC to 8 when the Trust state is changed to DSCP.
Receive Buffer	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added xon and xoff columns to the Receive Buffer configuration display.
4.2-1.0.0.0	
Physical Address Memory Allocation	[mlx5 Driver] Added support to register a specific physical address range.
Innova IPsec Adapter Cards	[Innova IPsec EN] Added support for NVIDIA Innova IPsec EN adapter card, that provides security acceleration for IPsec-enabled networks.

Category	Description
Precision Time Protocol (PTP)	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for PTP feature over PKEY interfaces. This feature allows for accurate synchronization between the distributed entities over the network. The synchronization is based on symmetric Round Trip Time (RTT) between the master and slave devices, and is enabled by default.
1PPS Time Synchronizat ion	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for One Pulse Per Second (1PPS) over IPoIB interfaces.
Virtual MAC	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for Virtual MAC feature, which allows users to add up to 4 virtual MACs (VMACs) per VF. All traffic that is destined to the VMAC will be forwarded to the relevant VF instead of PF. All traffic going out from the VF with source MAC equal to VMAC will go to the wire also when Spoof Check is enabled. For further information, please refer to "Virtual MAC" section in MLNX_OFED User Manual.
Receive Buffer	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added the option to change receive buffer size and cable length. Changing cable length will adjust the receive buffer's xon and xoff thresholds. For further information, please refer to "Receive Buffer" section in MLNX_OFED User Manual.
GRE Tunnel Offloads	 [ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for the following GRE tunnel offloads: TSO over GRE tunnels Checksum offloads over GRE tunnels RSS spread for GRE packets
NVMEoF	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for the host side (RDMA initiator) in RedHat 7.2 and above.
Dropless Receive Queue (RQ)	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for the driver to notify the FW when SW receive queues are overloaded.
PFC Storm Prevention	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for configuring PFC stall prevention in cases where the device unexpectedly becomes unresponsive for a long period of time. PFC stall prevention disables flow control mechanisms when the device is stalled for a period longer than the default pre-configured timeout. Users now have the ability to change the default timeout by moving to auto mode.

Category	Description
	For further information, please refer to "PFC Stall Prevention" section in MLNX_OFEDUser Manual.
Force DSCP	[ConnectX-4/ConnectX-4 Lx/ConnectX-5] Added support for this feature that enables setting a global traffic_class value for all RC QPs.
Q-in-Q	[ConnectX-5] Added support for Q-in-Q VST feature in ConnectX-5 adapter cards family.
Device Memory Programmin g [beta]	[ConnectX-5] Added support for on-chip memory allocation and usage in send/receive and RDMA operations at beta level.
Virtual Guest Tagging (VGT+)	[ConnectX-5] Added support for VGT+ in ConnectX-4/ConnectX-5 HCAs. This feature is s an advanced mode of Virtual Guest Tagging (VGT), in which a VF is allowed to tag its own packets as in VGT, but is still subject to an administrative VLAN trunk policy. The policy determines which VLAN IDs are allowed to be transmitted or received. The policy does not determine the user priority, which is left unchanged. For further information, please refer to "Virtual Guest Tagging (VGT+)" section in MLNX_OFED User Manual.
Tag Matching Offload	[ConnectX-5] Added support for hardware Tag Matching offload with Dynamically Connected Transport (DCT).
Shared Memory Region (MR)	 [ConnectX-3/ConnectX-3 Pro] Removed support for Shared MR feature on ConnectX-3/ConnectX-3 Pro adapter cards. As a result of this change, the following API/flags should not be used: ibv_exp_reg_shared_mr access shared flags for ibv_exp_reg_mr (IBV_EXP_ACCESS_SHARED_MR_XXX)
CR-DUMP	[All HCAs] Added support for the driver to take an automatic snapshot of the device's CR-Space in cases of critical failures. For further information, please refer to "CRDUMP" section in MLNX_OFED User Manual.
Upstream Libraries	[All HCAs] Added the option to install upstream libraries (based on upstream rdma-core) for DPDK users only. For further information, please refer to "Installing Upstream rdma-core Libraries" section in MLNX_OFED User Manual.
DiSNI	[All HCAs] Added the option to install libdisni package as part of MLNX_OFED.

Category	Description
	For further information, please refer to section "Installing libdisni Package" in MLNX_OFED User Manual.
Service Scripts	[All HCAs] Added the ability to disable the ' <i>stop</i> ' option in the openibd service script, by setting ALLOW_STOP=no in /etc/infiniband/openib.conf. Starting from the next release, ' <i>stop</i> ' option will be disabled by default, and in order to enable it, ALLOW_STOP should be set to 'yes' in the conf file, or <i>force-stop</i> should be run.
4.1-1.0.2.0	
RoCE Diagnostics and ECN Counters	[mlx5 Driver] Added support for additional RoCE diagnostics and ECN congestion counters under /sys/class/infiniband/mlx5_0/ports/1/hw_counters/ directory. For further information, refer to the <u>Understanding mlx5 Linux Counters</u> <u>and Status Parameters</u> Community post.
rx-fcs Offload (ethtool)	[mlx5 Driver] Added support for rx-fcs ethtool offload configuration. Normally, the FCS of the packet will be truncated by the ASIC hardware before sending it to the application socket buffer (skb). Ethtool allows to set the rx-fcs not to be truncated, but to pass it to the application for analysis. For more information and usage, refer to <u>Understanding ethtool rx-fcs for</u> <u>mlx5 Drivers</u> Community post.
DSCP Trust Mode	[mlx5 Driver] Added the option to enable PFC based on the DSCP value. Using this solution, VLAN headers will no longer be mandatory for use. For further information, refer to the <u>HowTo Configure Trust Mode on</u> <u>NVIDIA Adapters</u> Community post.
RoCE ECN Parameters	[mlx5 Driver] ECN parameters have been moved to the following directory: /sys/kernel/debug/mlx5/ <pci bus="">/cc_params/ For more information, refer to the <i>HowTo Configure DCQCN (RoCE CC)</i> <i>for ConnectX-4 (Linux)</i> Community post.</pci>
Flow Steering Dump Tool	[mlx5 Driver] Added support for mlx_fs_dump, which is a python tool that prints the steering rules in a readable manner.
Secure Firmware Updates	[mlx5 Driver] Firmware binaries embedded in MLNX_OFED package now support Secure Firmware Updates. This feature provides devices with the ability to verify digital signatures of new firmware binaries, in order to ensure that only officially approved versions are installed on the devices. For further information on this feature, refer to NVIDIA Firmware Tools (MFT) User Manual.

Category	Description		
Enhanced IPoIB	[mlx5 Driver] Added support for Enhanced IPoIB feature, which enables better utilization of features supported in ConnectX-4 adapter cards, by optimizing IPoIB data path and thus, reaching peak performance in both bandwidth and latency. Enhanced IPoIB is enabled by default.		
PeerDirect	[mlx5 Driver] Added the ability to open a device and create a context while giving PCI peer attributes such as name and ID. For further details, refer to the <u>PeerDirect Programming</u> Community post.		
Probed VFs	[mlx5 Driver] Added the ability to disable probed VFs on the hypervisor. For further information, see <u>HowTo Configure and Probe VFs on mlx5</u> <u>Drivers</u> Community post.		
Local Loopback	[mlx5 Driver] Improved performance by rendering Local loopback (unicast and multicast) disabled by mlx5 driver by default while local loopback is not in use. The mlx5 driver keeps track of the number of transport domains that are opened by user-space applications. If there is more than one user-space transport domain open, local loopback will automatically be enabled.		
1PPS Time Synchronizat ion (at alpha level)	[mlx5 Driver] Added support for One Pulse Per Second (1PPS), which is a time synchronization feature that allows the adapter to send or receive 1 pulse per second on a dedicated pin on the adapter card. For further information on this feature, refer to the <u>HowTo Test 1PPS on</u> <u>NVIDIA Adapters</u> Community post.		
Precision Time Protocol (PTP)	[mlx5 Driver] Added support for PTP feature in IPoIB offloaded devices. This feature allows for accurate synchronization between the distributed entities over the network. The synchronization is based on symmetric Round Trip Time (RTT) between the master and slave devices. The feature is enabled by default. For further information, refer to <u>Running Linux PTP with ConnectX-4</u> Community post.		
Fast Driver Unload	[mlx5 Driver] Added support for fast driver teardown in shutdown and kexec flows.		
HCAs: Connec	HCAs: ConnectX-5/ConnectX-5 Ex		
NVMEoF Target Offload	[ConnectX-5/ConnectX-5 Ex] Added support for NVMe over fabrics (NVMEoF) offload, an implementation of the new NVMEoF standard target (server) side in hardware. For further information on NVMEoF Target Offload, refer to <u>HowTo</u> <u>Configure NVMEoF Target Offload</u> .		

Category	Description
MPI Tag Matching	[ConnectX-5/ConnectX-5 Ex] Added support for offloading MPI tag matching to HCA.
RDMA CM	[All HCAs] Changed the default RoCE mode on which RDMA CM runs to RoCEv2 instead of RoCEv1. RDMA_CM session requires both the client and server sides to support the same RoCE mode. Otherwise, the client will fail to connect to the server. For further information, refer to <u>RDMA CM and RoCE Version Defaults</u> Community post.
Lustre	[All HCAs] Added support for Lustre file system open-source project.
4.0-2.0.2.0	
Operating Systems	Added support for Ubuntu v17.04.
4.0-2.0.0.1	
PCIe Error Counting	[ConnectX-4/ConnectX-4 Lx] Added the ability to expose physical layer statistical counters to ethtool.
Multiprotoco I Label Switching (MPLS) Tagged Packets Classificatio n	[ConnectX-4/ConnectX-4 Lx] Enabled packet flow steering rules with IPv4/IPv6 classification (for raw packet QP (DPDK) only) to work on IPv4/IPv6 over MPLS (Ethertype 0x8847 and 0x8848) encapsulated packets.
RoCE VFs	[ConnectX-4/ConnectX-4 Lx] Added the ability to enable/disable RoCE on VFs.
RoCE LAG	[ConnectX-4/ConnectX-4 Lx] Added support for RoCE over LAG interface.
Standard ethtool	[ConnectX-4/ConnectX-4 Lx] Added support for flow steering and rx-all mode.
SR-IOV Bandwidth Share for Ethernet/Ro CE (beta)	[ConnectX-4/ConnectX-4 Lx] Added the ability to guarantee the minimum rate of a certain VF in SR-IOV mode.
Adapter Cards	Added support for ConnectX-5 and ConnectX-5 Ex HCAs.

Category	Description
DSCP ConfigFS Control for RDMA-CM QPs	Added the ability to configure ToS/DSCP for RDMA-CM QPs only.
Soft RoCE (beta)	Add software implementation of RoCE that allows RoCE to run on any Ethernet network adapter whether it offers hardware acceleration or not.
NVMe over Fabrics (NVMEoF)	NVMEoF related module installation has been disabled by default. In order to enable it, add the " <i>with-nvmf</i> " installation option to the "mlnxofedinstall" script.
NFS over RDMA (NFSoRDMA)	Removed support for NFSoRDMA drivers. These drivers are no longer provided along with the MLNX_OFED package.

Customer Affecting Changes

Customer Affecting Change	Description
Customer	Affecting Changes 5.6-1.0.3.3
Interface Renamin g, PF/VF, Udev	 The OFED driver no longer performs Ethernet NetDev interface renaming for PFs and VFs. The udev rules file which implemented renaming (82-net-setup-link.rules) and its supporting script vf-net-link-name.sh are no longer installed by default. Renaming is thus performed by underlying mechanisms in udev, in the kernel, and in the BIOS. Users who wish to continue using the OFED driver renaming mechanism must add optioncopy-ifnames-udev to the OFED install command. To install these files at a later time, copy them from one of the following directories: /usr/share/doc/mlnx-ofa_kernel (RHEL8 and newer) /usr/share/doc/packages/mlnx-ofa_kernel (SLES) /usr/share/doc/mlnx-ofed-kernel-utils/examples (Debian-based releases)

Customer Affecting Change	Description
	 Note File 82-net-setup-link.rules should be copied to directory /etc/udev/rules.d File vf-net-link-name.sh should be copied to directory /etc/infiniband (make sure that it has both read and execute permission) After copying over the files, the driver should be restarted for the copied files to take effect Customers who wish prevent renaming of NetDev names should add "net.ifnames=0 biosdevname=0" to the kernel boot command line, and then reboot the host
Communi ty Operatin g Systems	Starting OFED 5.6, NVIDIA is introducing a new support model for OFED used on open source community operating systems. The goal of this new support model is to enable customers to use community-maintained variants of the Linux operating system, without being limited to major distributions that NVIDIA provides primary support for. For more information, see " <u>Installation on Community Operating Systems</u> " section in the user manual. For a list of supported Community OSs, please see " <u>Supported Community Operating Systems</u> " section in the release notes.
ar_mgr Subnet Manager Plugin	ar_mgr subnet manager plugin is no longer supported. For adaptive routing and SHIELD subnet manager configuration, please see the MLNX_OFED user manual.
	Starting UFM v6.7, Fabric Collector is no longer supported. For more information, see the <u>UFM release notes</u> .
Fabric Collector in UFM	(i) Note Please note that UFM is no longer bundled with OFED.

Customer Affecting Change	Description
OVS- DPDK— Partial Offload	Starting OFED 5.6, OVS-DPDK does not support partial offload.

Custome r Affecting Change	Description
5.5-1.0.3.2	
Disabling RoCE While Using sysfs	When using sysfs to enable/disable roce in kernel 5.5 and up, the "devlink reload" command (using iproute2 with devlink tool) will need to be used to activate the RoCE status change. Disable RoCE example: 1. echo 0 > /sys/bus/pci/devices/0000:08:00.0/roce_enable 2. devlink dev reload pci/0000:08:00.0
mlnx- ofa_kern el Installati on	 The source code for mlnx-ofa_kernel is no longer installed by default on RPM-based distributions (e.g., RHEL and SLES). Notes: mlnx-ofa_kernel is included in the << package mlnx-ofa_kernel-devel-source>> in the MLNX_OFED distributions under RPMS/ and may be manually installed from there. There is no change for deb-based distributions (Debian and Ubuntu). The full source is included, as before, in the package mlnx-ofed-kernel-dkms.
Software Encapsul ation Compati bility	There is an encapL2 compatibility issue with accelerated reformat action creation using mlx5dv_dr API. Using OFED 5.4 with firmware xx.32.1xxx and above or using OFED 5.5 with firmware lower than xx.32.1xxx will not allow accelerated reformat action. (Using OFED 5.4 and 5.5 with bundle firmware works properly.)
xpmem in RHEL8	Added xpmem packages in RHEL8 builds.
Python3	Starting OVS DPDK 2.15, the Python minimum required version is 3 and OVS- DPDK will not be compiled using Python 2.

Customer Affecting Change	Description	
Customer Affecting Changes 5.4-3.0.3.0		
CUDA, UCX, HCOLL	For UCX-CUDA and hcoll-cuda, CUDA was upgraded from version 10.2 to 11.2.	

Bug Fixes History

This table lists the bugs fixed in the last three major GA releases. For a list of old bug fixes, please refer to the release notes of the desired version.

Internal Reference Number	Description
	Description: Fixed an issue related to the driver's internal MR cache cleanup that caused high memory consumption on the host.
3946894	Keywords: MR cache cleanup
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.04-0.7.0.0
	Description: Fixed the setting of ATS for DMABUF MRs that caused some MRs to miss the ATS enablement. Lack of ATS enablement on DMABUF MRs results in slower performance when using these MRs.
3962426	Keywords: ATS, DMABUF, MRs
	Discovered in Release: 24.04-0.6.6.0
	Fixed in Release: 24.04-0.7.0.0
	Description: Fixed an issue related to the comparison process between the SW steering and FW steering modes to avoid kernel crashes incidences.
3798418	Keywords: SW steering, FW steering
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
3824463	Description: 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.

Internal Reference Number	Description
	Keywords: tc_wrap tool, VLAN
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
	Description: Fixed an IPSec issue that resulted in the disabling of Crypto offload.
3835880	Keywords: IPsec, Crypto offload
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
	Description: Fixed an issue that could have caused memory corruption when running XDP traffic.
3734180	Keywords: Memory, XDP traffic
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
	Description: Fixed an issue that caused IFA packets to be sent from the incorrect port on VF LAG.
3818623	Keywords: IFA packets, VF LAG
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
	Description: Added support for over KORG 6.7 read frequency_diff.
3866587	Keywords: Operating System, frequency
2000207	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
3851731	Description: 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.
	Keywords: LAG
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0

Internal Reference Number	Description
	Description: 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.
3546668	Keywords: PCI BAR size limitation
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.04-0.6.6.0
3764042	Description: Fixed an issue that resulted in PFC configuration being reset when updating the cable length via mlnx_qos.
	Keywords: PFC configuration, cable length
	Discovered in Release: 24.01-0.3.3.1
	Fixed in Release: 24.04-0.6.6.0

Internal Reference Number	Description
	Description: Fixed an issue that allowed concurrent creation of encap entries, and could potentially cause double free vulnerabilities.
3727823	Keywords: Encap entries, double free vulnerabilities
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.01-0.3.3.1
	Description: Resolved a discalculation issue where more Q-counters were freed than allocated when moving to switchdev mode.
3727026	Keywords: Q-counters, switchdev
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.01-0.3.3.1
3499413	Description: Fixed a kernel issue ("net/sched: act_ct: offload UDP NEW connections") where under heavy load, some connections could not be offloaded, leading to performance problems.
	Keywords: ASAP, CT offload

Internal Reference Number	Description
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.01-0.3.3.1
	Description: Fixed an issue that triggered a syndrome error due to debugfs entries that were not exposed for non supported general RoCE parameters, like rtt_resp_dscp.
3666162	Keywords: Syndrom error, debugfs entries
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.01-0.3.3.1
3710958	Description: Fixed an issue that triggered an error message due to outdated rule actions in the STE apply flow. The flow have been updated to check if the rule domain is different from the ASO CT action domain when applying the ASO CT action.
	Keywords: STE apply flow, ASO CT
	Discovered in Release: 23.10-1.1.9.0
	Fixed in Release: 24.01-0.3.3.1

Internal Reference Number	Description
	Description: Fixed an issue where an error was triggered in case devlink reload was attempted when there were allocated subfunctions.
3663363	Keywords: devlink reload, allocated subfunctions
	Discovered in Release: 23.10-0.5.5.0
	Fixed in Release: 23.10-1.1.9.0
	Description: Resolved an issue on ConnectX-4 Lx, where the VF state was not configured correctly following the activation of SR-IOV.
3660998	Keywords: ConnectX-4 Lx, VF state
	Discovered in Release: 23.10-0.5.5.0
	Fixed in Release: 23.10-1.1.9.0

Internal Reference Number	Description
	Description: Fixed an issue where changing the steering mode to firmware steering was unsupported for policy IPsec rules.
3653417	Keywords: Firmware steering
	Discovered in Release: 23.10-0.5.5.0
	Fixed in Release: 23.10-1.1.9.0

Internal Reference Number	Description
2612000	Description: Fixed an issue where the srp_daemon service was enabled by default. Starting from this release, the srp_daemon service is disabled by default, and can be enabled after MLNX_OFED installation using the systemctl enable srp_daemon command.
3613899	Keywords: srp_daemon
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
25 40 705	Description: Fixed an issue where the NVMe and mlx5_core drivers failed during BFB installation. As a result, Anolis OS could not be installed on the SSD and the mlxfwreset command did not work during Anolis BFB installation.
3549785	Keywords: NVMe, mlx5_core, Anolis OS
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue that occurred when a VF was set to get allmulti traffic. The issue caused the steering rules to send the multicast traffic received by the NIC back to the uplink.
3602955	Keywords: VF, allmulti traffic
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
3553766	Description: Fixed an issue where the enable_remote_dev_reset Devlink parameter was not supported on kernel versions below v5.10.

Internal Reference Number	Description
	Keywords: Devlink parameter
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where MAC address configuration for PFs could fail if SR-IOV was enabled at the same time.
3546694	Keywords: PF, MAC address, SR-IOV
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where firmware sync reset (with the 'mlxfwreset -d <device> -l 3 rsync 1' command) could fail on a system configured for hotplug on the PCIe slot on which the mlx5 card was mounted.</device>
3538018	Keywords: Firmware sync reset, mlx5 card
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where the enable_remote_dev_reset Devlink parameter was not supported on kernel versions below v5.10.
3587834	Keywords: Devlink parameter
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Resolved a warning that was triggered when starting the openibd service, which pertained to an unidentified 'ExecRestart' value within the 'Service' section.
3576351	Keywords: openibd, warning
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
3565980	Description: Fixed an issue where deb-based (Ubuntu and Debian) versions of MLNX_OFED did not install the rshim package, and the software needed to access the BlueField guest system from the host on a BlueField-type installation by default. Note that this was only

Internal Reference Number	Description
	needed (and enabled) on the host side. This issue did not occur on RPM-based systems (RHEL, SLES, Oracle, etc.).
	Keywords: Ubuntu, Debian, rshim package
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where the 'mlnx_tune -l' list of supported operating systems did not include several operating systems that were actually supported, such as RHEL8.6 and Ubuntu 22.04.
3557482	Keywords: mlnx_tune -l
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed a signature-related issue that occurred when installing DOCA on SLES15SP4 using the repository.
3549684	Keywords: DOCA, SLES15SP4
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where users who attempted to use OFED with Device ID NVD000000033, had to install the firmware manually.
3380263	Keywords: Device ID NVD000000033
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0
	Description: Fixed an issue where running rx-tls-offload over Korg6.0 as its TLS module did not work properly.
3228788	Keywords: NetDev, TLS
	Discovered in Release: 23.07-0.5.0.0
	Fixed in Release: 23.10-0.5.5.0

Internal Reference Number	Description
	Description: Resolved the kernel crash resulting from sysfs calls to profiles lacking TC (Traffic Control) support.
3546304	Keywords: sysfs calls, Trafic Control
	Discovered in Release: 23.04-0.5.3.3
	Fixed in Release: 23.07-0.5.0.0
	Description: Fixed an issue that prevented OS booting following an installation of the EN and RoCE drivers.
3531986	Keywords: OS booting, EN, RoCE
	Discovered in Release: 23.04-0.5.3.3
	Fixed in Release: 23.07-0.5.0.0
	Description: Fixed an issue in SLES 15 SP4 where the openibd service failed to start automatically after system boot.
3489233	Keywords: SLES 15 SP4,openibd, system boot
	Discovered in Release: 23.04-0.5.3.3
	Fixed in Release: 23.07-0.5.0.0
	Description: Fixed an issue that prevented the installation of OFED on RHEL systems using a non-default Python version.
3431430	Keywords: Installation, RHEL, Python
	Discovered in Release: 5.9-0.5.6
	Fixed in Release: 23.07-0.5.0.0
	Description: Fixed an OFED installation issue on BCLinux 21.10 that occurred when using the "add-kernel-support" installation flag.
3422823	Keywords: Installation, BCLinux 21.10, "add-kernel-support"
	Discovered in Release: 5.9-0.5.6
	Fixed in Release: 23.07-0.5.0.0
3264588	Description: Resolved a problem where the system boot process would hang when more than two Network Interface Cards were installed.
	Keywords: System boot, Network Interface Cards

Internal Reference Number	Description
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 23.07-0.5.0.0
	Description: Fixed an issue where the sysfs PHY counters displayed outdated information.
3499136	Keywords: sysfs PHY counters
	Discovered in Release: 23.04
	Fixed in Release: 23.07-0.5.0.0

Internal Reference Number	Description
	Description: Installing mlnx_tune on Python3 did not work properly. mlnx_tune now supports Python3 in addition to Python2.
2883451	Keywords: Installation, mlnx_tune, Python3
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 23.04-0.5.3.3
	Description: When creating a bond interface for all ports on a ConnectX-7 4-port HCA, the wrong bond name appeared in ibdev2netdev.
3219842	Keywords: RDMA, Bond Name, ibdev2netdev, ConnectX-7
	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 23.04-0.5.3.3
	Description: Changing traffic class via the sysfs while modifying QPs in parallel causes a deadlock.
3333919	Keywords: RDMA, TC, Sysfs, QP
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 23.04-0.5.3.3
3406019	Description: Due to a bug in the emulation layer, performance degradation might be experienced when running GPUDirect over Virtual Functions.

Internal Reference Number	Description
	Keywords: RDMA, GPUDirect, performance, VF
	Discovered in Release: 5.9-0.5.6.0
	Fixed in Release: 23.04-0.5.3.3
	Description: debugfs directories cannot be created for representors and sub-functions, thus the log might show error warning for either of the scenarios.
3233799	Keywords: NetDev, debugfs, SF, logging
	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 23.04-0.5.3.3
	Description: mlnx_tune script does not support Python3 interpreter.
1892663/18006	Keywords: mlnx_tune, Python3
33/2883451	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 23.04-0.5.3.3
3340542	Description: The verion number for perftest was not-standard resulting in some distribution packages recieving a higher version number than the OFED version for no good reason. Changed the naming of perftest to MAJOR.MINOR.PATCH.
	Keywords: Installation, perftest
	Fixed in Release: 23.04-0.5.3.3
	Description: knem did not fully support RHEL8.7 and newer releases.
2420775	Keywords: Installation, knem, RHEL
3428775	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 23.04-0.5.3.3
3431430	Description: Installing MLNX_OFED on a RHEL system that uses a non-default version of Python (e.g., Python3.9 on RHEL8.6, where the default is 3.6) may fail with an error that mlnx-tools is missing a dependency on 'python(abi)'. mlnx-tools includes a single script, mlnx_qos, that depends on a specific version of python. In such a case, after the fix, it may fail to run with such a non-standard version of Python.

Internal Reference Number	Description
	Keywords: Installation, Python, RHEL, mlnx-tools
	Discovered in Release: 5.9-0.5.6.0
	Fixed in Release: 23.04-0.5.3.3

Internal Reference Number	Description
3247519	Description: On an Ubuntu 22.04 system, when installing using the apt install method to install MLNX_OFED including Open vSwitch, and if the distribution Open vSwitch package was previously installed, the install may fail because of a left-over systemd generated file: the symbolic link /etc/systemd/system/openvswitch-switch.service.requires/ovs-record-hostname.service -> /lib/systemd/system/ovs-record-hostname.service .
	Keywords: Installation, Ubuntu 22.04, Open vSwitch
	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 5.9-0.5.6.0
	Description: Dapltest on RHEL9.x (ppc64le) could fail to run with a segmentation fault.
3296578	Keywords: Installation, RHEL9.x, Dapltest
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.9-0.5.6.0
3261289	Description: The host driver probe does not check whether there are existing SFs which are present in the device. As such, the host driver did not re-create those SFs.
0201200	Keywords: Core, Scalable Functions
	Fixed in Release: 5.9-0.5.6.0
3228719	Description: If there are multiple encapsulations and not all neighbors are valid, the kernel will go into panic mode.
	Keywords: ASAP ² , Kernel Panic
	Discovered in Release: 5.7-1.0.2.0

Internal Reference Number	Description
	Fixed in Release: 5.9-0.5.6.0
	Description: Moving to switchdev mode while deleting namespace may cause a deadlock.
2946873	Keywords: ASAP ² , Switchdev, Namespace
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.9-0.5.6.0
	Description: I n some topologies, like logical partitions, mlxfwreset is not supported.
3239291	Keywords: Core, mlxfwreset
	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 5.9-0.5.6.0
	Description: Creating external SFs on BF ARM when the host (x86) operating system does not support SFs may cause the host to crash.
3220855	Keywords: Core, Scalable Functions
	Discovered in Release: 5.8- 1.0.1.1
	Fixed in Release: 5.9-0.5.6.0

Internal Reference Number	Description
3253500	Description: The redundant freeing of a list item could lead to memory corruption, potentially causing the application to crash or incorrect traffic handling.
	Keywords: Steering, Memory Corruption, List, Pattern/Argument
	Fixed in Release: 5.8- 1.1.2.1
3214161	Description: The knem-dkms package explicitly requires GCC to build the knem driver (at install times). Under some circumstances, on Debian systems, the apt install method may result in a system that has only gcc- <version> (e.g., gcc-10) installed.</version>
	Keywords: Installation, Debian, GCC

Internal Reference Number	Description
	Fixed in Release: 5.8- 1.1.2.1
3230613	Description: Installing MLNX_OFED_LINUX on an Ubuntu system with CUDA (version < 11.6) may result in an automatic installation of the ucx-cuda package that will fail with an error message in the log file ucx-cuda.debinstall.log about missing dependencies.
	Keywords: Installation, Ubuntu, CUDA
	Fixed in Release: 5.8- 1.1.2.1
3235521	Description: The host driver probe did not check whether there are existing SFs which are present in the device, causing the host driver to not recreate those SFs.
	Keywords: Core, Scalable Functions
	Fixed in Release: 5.8- 1.1.2.1
	Description: If there are multiple encapsulations and not all neighbors are valid, the kernel will go into panic mode.
3228357	Keywords: ASAP ² , Encapsulation
	Discovered in Release: 5.5-1.0.3.2, 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.1.2.1
3232445	Description: When using BlueField with old kernels, multiple OVS meter do not work.
	Keywords: ASAP ² , BlueField, Meter, OVS, Offload
	Fixed in Release: 5.8- 1.1.2.1

Internal Reference Number	Description
3234066	Description: When configuring IPsec full offload, after sending traffic for approximately 30 minutes, the traffic stops at some point and the connection gets lost.
	Keywords: Steering, SMFS, Matcher Disconnect
	Fixed in Release: 5.8- 1.0.1.1

Internal Reference Number	Description
3179535	Description: SMFS will try to merge flow rules with the same matching criteria (as they share the same matcher) into one multi-destination rule. If merging fails, the matcher is disconnected by mistake.
5115555	Keywords: Steering, SMFS, Matcher Disconnect
	Fixed in Release: 5.8- 1.0.1.1
	Description: ibv_reg_mr for huge pages was optimized in kernel >= 5.12
3214198	Keywords: RDMA, ibv_reg_mr
3214198	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
	Description: Moving to SwitchDev mode while deleting namespace over Linux-6.0 can sometimes cause a deadlock.
2984134	Keywords: RDMA, SwitchDev
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.8- 1.0.1.1
	Description: A net device validation issue prevented running IPv6 traffic using an RDMA communication manager between two interfaces on same host with same subnet.
3106228	Keywords: RDMA, IPv6, Communication Manager
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.8- 1.0.1.1
	Description: I n mlx5dv_mkey_check manpage, there is an i naccurate description of signature error handling flow.
3151843	Keywords: RDMA, manpage
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
3229002	Description: Creating and deleting MRs, caused a kernel slab cache leak issue.
	Keywords: RDMA, Cache
	Discovered in Release: 5.7-1.0.2.0

Internal Reference Number	Description
	Fixed in Release: 5.8- 1.0.1.1
	Description: The rdma res show cm_id command does not list all cm_ids when some of them are in LISTEN state.
3236217	Keywords: RDMA, cm_ids
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.8- 1.0.1.1
	Description: In older kernel version, PTP was not supported over VLAN interfaces.
3146128	Keywords: NetDev, PTP, VLAN
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
	Description: HW-GRO feature was blocked due to firmware limitations.
2000772	Keywords: NetDev, HW-GRO
2969772	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.8- 1.0.1.1
	Description: STP packets failed to be transmitted.
2006202	Keywords: NetDev, STP
3096393	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.8- 1.0.1.1
	Description: When using sysfs to read the hash function used to distribute the traffic between the T I R s (Transpo rt Interface Receive) , on occasion, the server crashed.
3236984	Keywords: NetDev, sysfs
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
3126000	Description: Upgrading from version 5.6-2 to 5.7 failed.
	Keywords: Installation
	Discovered in Release: 5.6-2.0.9.0

Internal Reference Number	Description
	Fixed in Release: 5.8- 1.0.1.1
	Description: Building with KMP enabled fails due to missing packages. OFED packages will now be built with KMP disabled.
3230524	Keywords: Installation, KMP
	Fixed in Release: 5.8- 1.0.1.1
	Description: The script <u>install.pl</u> , used for (re)building kernel modules, used the name "kernel-source" as the package of the kernel-source on SLES systems.
3158725	Keywords: Installation, SLES
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.8- 1.0.1.1
3142212	<pre>Description: Starting firmware version xx.34.0350, a new NVCONFIG has been added to the ARM side only: MANAGEMENT_PF_MODE. If this config is on, the user will see a PCI Function (PF) which failed to probe: [6.837102] mlx5_core 0000:03:00.2: mlx5_cmd_check:756: (pid 206): ENABLE_HCA(0x104) op_mod(0x0) failed, status bad parameter(0x3), syndrome (0x6ca1f5) [6.864227] mlx5_core 0000:03:00.2: mlx5_peer_pf_init:40: (pid 206): Failed to enable peer PF HCA err(-22) [6.883453] mlx5_core 0000:03:00.2: mlx5_load:1129:(pid 206): Failed to init embedded CPU [8.261268] mlx5_core 0000:03:00.2: init_one:1365:(pid 206): mlx5_load_one failed with error code -22 [8.280056] mlx5_core: probe of 0000:03:00.2 failed with error -22</pre>

Internal Reference Number	Description
	Keywords: Installation
	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
	Description: Using a 1-CPU system casues possible command flush deadlock.
3174928	Keywords: Core
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.8- 1.0.1.1
	Description: An incorrect termination table was used with the uplink-to-uplink forward rule.
3228721/ 3228357	Keywords: ASAP ² , eSwitch
3228357	Discovered in Release: 5.7-1.0.2.0
	Fixed in Release: 5.8- 1.0.1.1
3220120	Description: In old kernels, when a VXLAN tunnel is set up on one OVS bridge and PF is up on another OVS bridge, traffic does not offload as expected.
	Keywords: ASAP ² , VXLAN
	Discovered in Release: 5.4-3.0.3.0
	Fixed in Release: 5.8- 1.0.1.1

Internal Reference Number	Description
3032335	Description: Creating multiple steering rules that modify a packet and match on the same packet headers can cause an error to be displayed in dmesg when deleting the steering rules.
	Keywords: Steering Rules
	Fixed in Release: 5.7-1.0.2.0
3011368	Description: Some IB spec QP state behaviour on post_send()/recv() is not being fully enforced. The fix makes the QP complaint to IB spec

Internal Reference Number	Description
	about when it is allowed to post_send()/recv() and when it should return an error.
	Keywords: RDMA, IB spec QP
	Fixed in Release: 5.7-1.0.2.0
	Description: When changing trust state from PCP to DSCP, the TC number changes by default to 8, in some cases, disrupting traffic prioritization if trust state is changed back to PCP.
3075125	Keywords: NetDev, QoS
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.7-1.0.2.0
	Description: In the current release, the following OPNs/PSIDs should be manually upgraded: MCX753106AS-HEA-N NVD000000023 MCX75310AAS-HEA-N NVD000000024
3054413	Keywords: ConnectX-7, Upgrade
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.7-1.0.2.0
	Description: In versions of MLNX_OFED before 5.7, the xpmem kernel module was not signed. When it was installed on systems (mostly RHEL and other compatible systems) the following error message would appear: "xpmem: loading out-of-tree module taints kernel."
3070653	Keywords: Installation, xpmem
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.7-1.0.2.0
3075357	Description: In Debian-based distributions, in /etc/init.d/openibd, the path to enable the firmware tracer is /sys/kernel/debug/tracing/events/mlx5/fw_tracer/enable instead of /sys/kernel/debug/tracing/events/mlx5/mlx5_fw/enable . As a result, firmware tracer will never get enabled even when supported.
	Keywords: Installation, Kernel Trace Debug
	Discovered in Release: 5.6-1.0.3.3

Internal Reference Number	Description
	Fixed in Release: 5.7-1.0.2.0
	Description: The minimum Tx rate limit is not supported with link speed of 1Gb/s.
2688191	Keywords: Rate Limit, 1Gb/s
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.7-1.0.2.0
	Description: Destroying mlxdevm group while SF is attached to it is not supported.
3044255	Keywords: ASAP ² , mlxdevm, QoS, Group, Scalable Functions
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.7-1.0.2.0
	Description: Using OVS offload with NIC mode (non switchdev mode) causes traffic to drop.
3047142	Keywords: ASAP ² , Offload, NIC Mode, OVS
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.7-1.0.2.0
	Description: In some cases VF metering configuration failure caused a deadlock.
3123986	Keywords: ASAP ² , VF Metering
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.7-1.0.2.0
	Description: A race condition may cause some connection aging to set to 24 hours instead of 30 seconds.
3053842	Keywords: ASAP ² , Connection Tracking, Aging
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.7-1.0.2.0

Internal Reference Number	Description
	Description: When an already-loaded 'non-mellanox' auxiliary device on the auxiliary bus OFED driver exists, load may fail and cause kernel panic.
3079038	Keywords: Driver Load
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.6-2.0.9.0
2000222	Description: On SLES15 systems that have both python3 and python2 installed, rebuilding kernel modules fails with an error in the mlnx-tools package, and specifically in the mlnx-tools build log, about missing ib2ibsetup.8.
3066233	Keywords: Installation
	Discovered in Release: 5.6-1.0.3.3
	Fixed in Release: 5.6-2.0.9.0

Internal Reference Number	Description
	Description: NULL dereference may occur when performing port up/down.
2858237	Keywords: Host
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.6-1.0.3.3
	Description: Reloading devlink in NetDev profile caused deadlock.
2607442	Keywords: Devlink
2697443	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.6-1.0.3.3
2947734	Description: When using NFS over RDMA rpcrdma.ko created some entry files under /proc folder (e.g., "-rw-rr 1 root root 0).
	Keywords: NFS, RDMA
	Discovered in Release: 5.5-1.0.3.2

Internal Reference Number	Description
	Fixed in Release: 5.6-1.0.3.3
	Description: ECPF port was not properly recognized.
2905055	Keywords: ECPF
	Fixed in Release: 5.6-1.0.3.3
	Description: Gratuitous ARP during rdma_connect is not handled properly.
2771739	Keywords: Gratuitous ARP
	Fixed in Release: 5.6-1.0.3.3
	Description: A locking issue in steering rules deletion, at times, could cause a deadlock while inserting or deleting new rules.
2888178	Keywords: Deadlock, Steering
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: ibdump sources were not shared under the sources directory.
2979734	Keywords: ibdump Sources
	Fixed in Release: 5.6-1.0.3.3
	Description: Crypto offload of UDP traffic on top of IPv6 was unsupported.
2820245	Keywords: IPsec, Crypto, Offload
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2869109	Description: IPsec crypto offload for non TCP/UDP encapsulated traffic broke.
	Keywords: IPsec, Crypto, Offload
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2905896	Description: Leaving a multicast group (rdma_leave_multicast) used the wrong address and left the interface in the multicast group.

Internal Reference Number	Description
	Keywords: RoCE, Multicast
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2939691	Description: Unsupported parameters were ignored. Now, when using unsupported syntax or unsupported command line parameters, the application will fail with an error message.
200001	Keywords: Command Line, Parameters
	Fixed in Release: 5.6-1.0.3.3
3024670	Description: When many MRs are allocated, the driver searches for a free MR with the best size fit. When the cache is dry, a huge MRs (1 MB) instead of small once (32B) could be selected because there was no better fit. The memory overhead was limited to avoid synchronous MR creation.
	Keywords: Memory, MR
	Fixed in Release: 5.6-1.0.3.3
	Description: In the rdma-core library, the CMA device was retrieved in the wrong way when libnl is not used.
3020746	Keywords: libnl, rdma-core
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.6-1.0.3.3
	Description: Ethtool that is part of the original OFED package failed to dump correct EEPROM values when using -m flag.
2939037	Keywords: Ethtool, EEPROM
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: Manual pages were missing from infiniband-diags tools. Added manpages for all OSs that provide the python-docutils package.
2582616	Keywords: Manual Pages (manpages), infiniband-diags
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.6-1.0.3.3

Internal Reference Number	Description
	Description: An increment of count variable was missing when looping over output buffer in mlx5e_self_test(). As a result, the garbage value of ethtool -t was resolved.
2979137	Keywords: ethtool, selftest
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2752622	Description: On SLES 15, the inbox modules in the directory mlxsw (such as mlxsw_spectrum) was not supported. When they were installed when installing MLNX_OFED, they no longer worked (as they depend on a different version of the mlx* modules) and could cause an error at time of installation.
	Keywords: Installation
	Discovered in Release: 5.4-3.0.3.0
	Fixed in Release: 5.6-1.0.3.3
	Description: In cases where MLNX_OFED was reinstalled on a certain system without usingforce, the installation could fail requiring the removal of the infiniband-diags package.
2200320	Keywords: Installation, Force, infiniband-diags
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.6-1.0.3.3
	Description: When uninstalling the kmod-xpmem package, xpmem module was not unloaded. From now on, after uninstalling, xpmem module will be removed automatically.
2984013	Keywords: Installation, xpmem
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2984098	Description: OFED installation modified file "/etc/yum.conf" to exclude some packages from the Yum repositories. As of RHEL 8, /etc/yum.conf is a symlink to /etc/dnf.conf and this edit breaks the symlink. As there is no use in such an edit, OFED no longer edits this file.
	Keywords: Installation, Yum Repositories, RHEL

Internal Reference Number	Description
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: PCIe device address to RDMA device name mapping on x86 host may change after the driver restarts in Arm.
2934828	Keywords: RDMA, Arm, Driver
	Fixed in Release: 5.6-1.0.3.3
	Description: In some cases, the firmware tracer did not work with NEO-Host.
2946450	Keywords: NEO-Host, Firmware Tracer
	Fixed in Release: 5.6-1.0.3.3
	Description: current_link_speed sysfs was missing.
0047045	Keywords: sysfs
2947645	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: If the commands were not entered in the correct order when setting buffer size and allocation using the mlnx_qos command, on some occasions, the xoff_threshold calculation broke pausing functionality.
3025582	Keywords: Driver, xoff, Buffer
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: First generation BlueField SoC based DPUs were not supported.
2859206	Keywords: BlueField, SoC
	Fixed in Release: 5.6-1.0.3.3
2936867	Description: Creating a TC rules with more than 30 actions caused kernel panic.
	Keywords: ASAP ² , Call Trace, TC
	Discovered in Release: 5.5-1.0.3.2

Internal Reference Number	Description
	Fixed in Release: 5.6-1.0.3.3
	Description: IP-in-IP packets received in one queue instead of hashing to multi queues.
3016685	Keywords: NetDev, Tunneling, RSS
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.6-1.0.3.3
	Description: Fixed compatibility issue of mlnx_qos for python3.9 deprecated tostring/fromstring.
3023304	Keywords: Python3, Compatibility
	Fixed in Release: 5.6-1.0.3.3
2007207	Description: IPsec flow tables design caused the number of IPsec tunnels to be limited to 16K. Changed the flow tables design to support up to 32K IPsec tunnels per protocol (IPv4/IPv6).
2887387	Keywords: IPsec
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.6-1.0.3.3
	Description: When configuring over 1000 IPsec sessions caused performance issues.
2887394/2887	Keywords: IPsec
381	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: Hairpin rules failed to send packet back to wire when IPsec full offload is enabled.
2906002	Keywords: IPsec Full Offload, Hairpin
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2901506	Description: On rare occasions, the application did not use any raw WQE feature and unexpectedly got wc opcode IBV_WC_DRIVER2.

Internal Reference Number	Description
	Keywords: RDMA, Raw WQE, IB_WC_DRIVER2
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
	Description: Under certain conditions, incorrect handling of resources caused memory corruption over software steering resources leading to failure of OVS to offloaded the traffic to the hardware.
2890024	Keywords: ASAP ² , Steering, OVS, Memory
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2874200	Description: Using hairpin tunnel traffic, caused incorrect TC rules to be created. Example: tunnel(tun_id=0×65,src=10.10.11.3,dst=10.10.11.2,ttl=0/0,tp_dst=4789,flags(+key)),,in_port (vxlan_sys_4789),, actions:set(tunnel(tun_id=0×66,src=10.10.12.2,dst=10.10.12.3,tp_dst=4789,flags(key))),vxlan _sys_4789
	Keywords: ASAP ² , Hairpin, OVS, SwitchDev
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3
2891499	Description: Adding a route with next hop object caused a warning in dmesg and could possibly lead to kernel panic.
	Keywords: ASAP ² , Route, SwitchDev, Call Trace, Nexthop
	Discovered in Release: 5.5-1.0.3.2
	Fixed in Release: 5.6-1.0.3.3

Internal Reference Number	Description
2842077	Description: Between scripts there was a possibility for Inconsistency in python3 header line (shebang line) because some distributions may no longer have /usr/bin/python.

Internal Reference Number	Description
	Keywords: Python3
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.5-1.0.3.2
	Description: The driver didn't set the PCP-based priority for DCT, hence DCT response packets were transmitted without user priority.
2792432	Keywords: User Priority, DCT
	Fixed in Release: 5.5-1.0.3.2
	Description: Node reboots may trigger memory corruption in OFED CM.
2424200	Keywords: CM Migration
2434399	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.5-1.0.3.2
2696789	Description: Redesigned the locks around peer MR invalidation flow to avoid a potential deadlock as Peer-direct patch may cause deadlock due to lock inversion. Notes:
	 For GPU drivers prior to r470, the user should update nv_peer_mem to the next version, probably 1.2. For GPU drivers from r470 or later branches shipped with nvidia-peermem, the driver will have an option to update to newer releases which take advantage of the redesigned MLNX_OFED support.
	Keywords: lock inversion, nv_peer_mem
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.5-1.0.3.2
2792480	Description: Running tcpdump on bonding standby port caused to lose the network.
	Keywords: NetDev
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.5-1.0.3.2

Internal Reference Number	Description
	Description: Running yum update will upgrade kylin-release to a higher version. The version of this package is used for kylin10sp2 detection so the script will detect kylin 10 instead of kylin10sp2 and use its repository by mistake.
2782406	Workaround: Upgrade, kylin
	Discovered in Release: 5.4-3.0.3.0
	Fixed in Release: 5.5-1.0.3.2
	Description: Starting from GPU Driver version r465, nv_peer_mem was shipping in the GPU driver package under the name nvidia-peermem. Updating OFED required nvidia-peermem rebuild, otherwise it was stubbed out by the kernel.
2736003	Keywords: Installation, GPU Driver
	Discovered in Release: 5.4-3.0.3.0
	Fixed in Release: 5.5-1.0.3.2
	Description: xpmem driver is not supported on PowerPC.
2823700	Keywords: Installation, xpmem, PowerPC
	Fixed in Release: 5.5-1.0.3.2
	Description: Suspend flow freed the VLAN data so the data was not restored during the resume flow.
2802508	Keywords: VLAN, Suspend Flow, Resume Flow
	Fixed in Release: 5.5-1.0.3.2
	Description: Connection tracking rules with fragmentation had 0 stats.
2796010	Keywords: BlueField, Connection Tracking, Fragments, ASAP ²
2790010	Discovered in Release: 5.4-2.4.1.3
	Fixed in Release: 5.5-1.0.3.2
2803403	Description: Traffic failed to pass when OVS bridge is configured with bond interface and IP is configured over the OVS internal (bridge) port.
	Keywords: Bond, VF LAG, OVS, Internal Port, ASAP ²
	Discovered in Release: 5.2-1.0.4.0

Internal Reference Number	Description
	Fixed in Release: 5.5-1.0.3.2
	Description: VXLAN with IPsec crypto offload does not work.
2420202	Keywords: VXLAN; IPsec crypto
2438392	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.5-1.0.3.2
	Description: Conducting a driver restart while in VF LAG mode may cause unwanted behaviour such as kernel crashes.
2677225	Keywords: ASAP ² , Bonding, Driver Restart, VF LAG
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.5-1.0.3.2

Internal Reference Number	Description
	Description: In version 5.4, there was some offload breakage when using OVS.
2852904	Keywords: TSO, UDP Tunnels
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.1.0.0
	Description: Running tcpdump on a bonding standby port resulted in the loss of the network.
2792480	Keywords: NetDev
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
2696789	Description: Redesigned the locks around peer MR invalidation flow to avoid a potential deadlock as Peer-direct patch may cause deadlock due to lock inversion. Notes:
	• For GPU drivers prior to r470, the user should update nv_peer_mem to the next version, probably 1.2.

Internal Reference Number	Description
	• For GPU drivers from r470 or later branches shipped with nvidia- peermem, the driver will have an option to update to newer releases which take advantage of the redesigned MLNX_OFED support.
	Keywords: lock inversion, nv_peer_mem
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: A race that resulted in a QCE with an error, caused errors in UMR QP. To prevent the UMR QP from getting into error, we fixed the MR deregistration flow (e.g., Peer Ikey which is always revoked before destroying it).
2739689	Keywords: QCE, UMR
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
2691656	Description: When using bonding, ibdev2netdev would sometimes match the infiniband device to the net device bonding interface, and sometimes to the underlying Infiniband net device interface. ibdev2netdev now skips InfiniBand net device bonding interfaces, and always matches InfiniBand devices to the underlying InfiniBand net device interfaces.
	Keywords: ibdev2netdev Bonding
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.4-3.0.3.0
	Description: Fixed Decap flows inner IP_ECN match to take into account software modification of the match value according to RFC 6040 4.2.
2687643	Keywords: decap, ASAP ² , ECN, RoCE
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
2691081	Description: Removed metadata from the rpm package mlnx-ofa_kernel where it claimed to Provide an older version of rdma-core. This made sense in older versions where we needed to avoid installing rdma-core.

Internal Reference Number	Description
	But does not make sense anymore. And caused problems to some users installing rdma-core-devel through meta-packages.
	Keywords: Installation
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: Removed manual build-time file list generation in mlnx-tools. Only keep it for python-installed files. And avoid guessing the version of python we use and the directory to which we install.
2727062	Keywords: Installation
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: Removed useless build-time editing of uninstall.sh in ofed- scripts that caused the build to fail (in the case ofadd-kernel-support) in some rare cases.
2708220	Keywords: Installation
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: Some Dell OFED Factory Installation packages were missing dependencies. Removed the package rdma-core-devel from the Dell MLNX_OFED packages as it was not needed and some of its dependencies are not included.
2730547	Keywords: Installation, Dell
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
2699662	Description: MLNX_OFED build scripts fixed to also build hcoll with CUDA support on RHEL8 x86_64 platforms.
	Keywords: Installation, CUDA
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0

Internal Reference Number	Description
2686877	Description: Changing mtu takes too long. Reduced number of calls to synchronize_net to once for all channels.
	Keywords: mtu, synchronize_net
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
2748328	Description: When trying to upgrade a kmp package, it conflicts and needs user help to choose whether to replace it or not. The fix avoids conflicts from /usr/lib/rpm/kernel-module-subpackage script which was changed in the builder. Building the packages with kmp enabled on the other image will cause the issue to reproduce.
	Keywords: Upgrade, kmp Package
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: On Ubuntu and Debian systems for openvswitch-switch (in case installing using e.govs-dpdk orwith-openvswitch), the installer misses a run-time dependency of libpcap0.8.
2707023	Keywords: Installation, Ubuntu, Debian
	Discovered in Release: 5.4-1.0.3.0
	Fixed in Release: 5.4-3.0.3.0
	Description: The full path to the directory that contains the installer must not contain a space or any similar white-space character, otherwise the installer will fail.
2563366	Keywords: Installation, White Space
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-3.0.3.0

Internal Reference Number	Description
2684302	Description: To support scalability, function representor channels were limited to 4. However in scenarios when SF are not used, certain use cases require representors to support a large number of channels.

Internal Reference Number	Description
	Hence, representor channel limit to 4 is applicable only when a PCI device, such as Scalable Function support, is enabled.
	Keywords: Representor Channels
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
	Description: Matching on ipv4_ihl (internet header length) was supported only for outer headers. Support has been added for inner headers too.
2644217	Keywords: Internet Header Length, ipv4_ihl
	Fixed in Release: 5.4-1.0.3.0
2626906	Description: When using one counter for both pop/push VLAN actions, the counter value is incorrect. Split the counter for pop_vlan_action_counter and push_vlan_action_counter.
2020900	Keywords: Pop/Push VLAN
	Fixed in Release: 5.4-1.0.3.0
2653382	Description: Incorrect L3 decapsulation occurs when the original inner frame is small and was padded to comply with minimum frame size of 64-bytes.
2033302	Keywords: SW Steering, Decapsulation, Padding
	Fixed in Release: 5.4-1.0.3.0
	Description: dapl and libmlx4 are needed by libdat2 and libdpdk. In order to remove or update dapl, its dependencies need to be removed.
2612725	Keywords: dapl, libmlx4
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.4-1.0.3.0
	Description: An override of log_max_qp by other devices occurs if the devices share the same mlx5_core module.
2649134	Keywords: log_max_qp, mlx5_core
	Fixed in Release: 5.4-1.0.3.0

Internal Reference Number	Description
2638029	Description: A synchronization issue where closing and opening channels (which may happen on configuration changes such as changing number of channels) may cause null pointer dereference in function mlx5e_select_queue.
	Keywords: mlx5e_select_queue, Synchronization, Tx
	Fixed in Release: 5.4-1.0.3.0
	Description: Enabling tx-udp_tnl-csum-segmentation has no effect on the driver. tx-udp_tnl-csum-segmentation has been moved to "off [fixed]".
2678982	Keywords: tx-udp_tnl-csum-segmentation
	Discovered in Release: 5.4-0.5.1.1
	Fixed in Release: 5.4-1.0.3.0
	Description: Some MLNX_OFED dkms packages ignored (install-time) build errors and considered the packages properly built. Those errors are now not ignored and indicated as package installation errors.
2610870	Keywords: dkms, Installation
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.4-1.0.3.0
2617820	Description: Old udevd versions could get stuck renaming network devices, leaving interfaces named eth* instead of enp*. Updating the systemd version resolves this issue. For example, if an issue detected on RHEL 7.6 with systemd-219-62, updating the systemd version to systemd-219-67 resolves the issue.
	Keywords: udev, systemd, RHEL
	Discovered in Release: 5.4-0.5.1.1
	Fixed in Release: 5.4-1.0.3.0
	Description: Flows with t commit action with ct state -trk are not be offloaded (i.e., table=0,ct_state=-trk,ip actions=ct(commit,table=1)).
2632768	Keywords: ASAP ² , Connection Tracking
	Fixed in Release: 5.4-1.0.3.0

Internal Reference Number	Description
	Description: Connection tracking over VF LAG with tunnel encapsulation/decapsulation is not supported and may cause traffic drop.
2247143	Keywords: ASAP ² , Connection Tracking, VF LAG, Tunnel
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.4-1.0.3.0
	Description: When stack size is limit to 1024, OFED compilation fails.
2507227	Keywords: Compilation, Stack
2597327	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
	Description: Setting rate/burst values higher than 2,147,483,648 are rejected.
2609641	Keywords: VF Metering
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
	Description: Offloaded remote mirroring flows on tunnel device caused forwarded traffic to VF to not be decapsulated.
2626920	Keywords: ASAP ² , Offload, Remote Mirroring, Tunnel
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
2660247	Description: Trying to set VPort match mode on VF (cat/sys/class/net/enp8s0f2/compat/devlink/vport_match_mode), leads to kernel crash.
	Keywords: ASAP ² , Kernel Crash
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
2667484	Description: OVS flows are not being offloaded over socket-direct devices.

Internal Reference Number	Description
	Keywords: ASAP ² , Socket-Direct
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
	Description: When VXLAN is configured and illegal route is added, the system crashes with call trace.
2663042	Keywords: ASAP ² , Offload, Tunnel, Call Trace
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
	Description: If any traffic is sent before the netdev goes up for the first time, a division by zero caused by a modulo operation may occur in ndo_select_queue, leading to a kernel panic.
2354761	Keywords: NetDev; ndo_select_queue
	Discovered in Release: 5.3-1.0.0.1
	Fixed in Release: 5.4-1.0.3.0
2562053/266 7551	Description: After restarting driver, the x86 host may be in grace period and may not recover on its own. As part of the fix, 5 FW_fatal recoveries are allowed within the 20-minute grace period. As a result, the grace period in the devlink health show command will appear as 0 for FW_fatal reporter.
	Keywords: BlueField Reload, recovery, reset flow
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.4-1.0.3.0

Internal Reference Number	Description
2635638	Description: In fork situation, if parent/children processes happen to have same virtual address, then the doorbell mechanism may not work well and may leads to errors in application behavior.
	Keywords: RDMA, Fork

Internal Reference Number	Description
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.4-1.0.3.0

Internal Reference Number	Description
2393352	Description: Using "with-openvswitch" flag during MLNX_OFED installation may not work on Debian 10 systems.
	Keywords:with-openvswitch, Debian
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: ib_uverbs module parameter disable_raw_qp_enforcement is deprecated and should no longer be used.
2445058	Keywords: disable_raw_qp_enforcement, ib_uverbs
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue in ConnectX-5 and earlier that when the module is missing, the driver reported a connector type that is different than OTHER.
2434650	Keywords: Module, Connector Type
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Solved a compilation error by fixing a backport issue with unpin_user_pages_dirty_lock function.
2434650	Keywords: Memory, unpin_user_pages_dirty_lock
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
2434650	Description: Fixed an issue that caused a deadlock on IPoIB interface. While IPoIB driver handes multicast groups, the fix assures that

Internal Reference Number	Description
	everything is done under safe lock while handled.
	Keywords: IPoIB, Multicast
	Discovered in Release: 4.6
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where VLAN header was not popped on VF Rx when the eSwitch priority tagging was configured.
2505615	Keywords: ASAP ² , Priority Tagging, VLAN
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed connection tracking (CT) offload in NIC mode by using correct steering domain for the rules.
2494257	Keywords: ASAP ² , Connection Tracking, NIC Mode
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where offload of rules from OVS internal port to uplink failed.
2461213	Keywords: ASAP ² , OVS
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue in the tunnel mishandling that can happen when the tunnel overlay device is an OVS internal port.
2444523	Keywords: ASAP ² , OVS internal port offloading
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
2566354	Description: Fixed incorrect parsing of network configuration when the optionnet (-n) was given to mlnxofedinstall: get network configuration from the output of 'ip' instead of 'ifconfig'.
	Keywords: Installation

Internal Reference Number	Description
	Fixed in Release: 5.3-1.0.0.1
	Description: Dropped unsupported devices from OFED rdma-core description.
2495065	Keywords: rdma-core
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Backported MLNX_OFED kernel to support elrepo 5.8 kernel.
2482696	Keywords: add-kernel-support
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed ability to build xpmem on kernel version 5.6.
2401104	Keywords: add-kernel-support, xpmem
2481104	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where kernel build on SLES 15 systems that configures scripts assume SLES 15 systems have /etc/SuSE-release or /etc/SUSE-brand. These files no longer exist on SLES 15.
2440062	Keywords: add-kernel-support, SLES 15
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where running data on Geneve tunnel on a VF may result in CQE error and a failure t to transmit data.
2445146	Keywords: Virtual Function
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
2494008	Description: Fixed an issue where the driver silently ignores the settings of an already-set ECN value (0->0, 1->1) via sysfs.
	Keywords: RDMA, ECN

Internal Reference Number	Description
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where KVS offload, under certain conditions, takes too long. Improved malloc performance by increasing the memory reuse and reducing the stress on malloc and free.
2581127	Keywords: MLNX5DR, Software Steering
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
2502564	Description: Fixed an issue where when using switchdev mode with SMFS, inserting duplicate rules from userspace was not supported (required when there are a few instances of the same application). As part of the fix, added support for update_fte which is called in case a duplicate rule is being added.
	Keywords: SwitchDev, Steering
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where creating 127 ports on each VF may fail as the current kernel does not support an RDMA device with more than 255 ports.
2433351	Keywords: VF, RDMA, virtualization
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where changing the "other" channels count by "ethtool -L other " command on Kernel 5.10 may cause a kernel panic.
2333971	Keywords: Kernel 5.10, kernel panic, ethtool, "other" channels
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
2454952	Description: Fixed an issue where MLNX_OFED cannot be built on top of Kernel 5.4.87.
	Workaround: operating system, kernel

Internal Reference Number	Description
	Discovered in Release: 5.2-2.2.0.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Fixed an issue where Switch and eSwitch offloads are not supported for SR-IOV and its sub functions when installing MLNX_OFED over upstream kernel v5.10 or higher.
2383355	Keywords: eSwitch, Kernel, SR-IOV
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.3-1.0.0.1
	Description: Creating a bond via NetworkManager and restarting the driver (openibd restart) results in no pf0hpf and bond creation failure.
2278833	Keywords: Bond, LAG, network manager, driver reload
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-2.2.0.0
	Description: In rare cases, mlx5_cmd_exec may lead to recount warning, such as: WARNING: CPU: 1 PID: 30811 at lib/refcount.c:28 refcount_warn_saturate+0xd9/0xe0 This refcount warning can be ignored.
2293460	Keywords: mlx5_cmd_exec
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Prior to stopping Open vSwitch, the following command should be run. ovs-appctl exitcleanup
2328206	Keywords: Open vSwitch, ovs-appctl
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.2-2.2.0.0
2278352	Description: Under Containers environment, having "rdma-core" package installed on the Hypervisor may conflict with loading updated MLNX_OFED drivers from within a Container (using "/etc/init.d/openibd restart" for example). The "rdma-core" services may load the Inbox drivers on the Hypervisor, causing the loading of the MLNX_OFED

Internal Reference Number	Description
	drivers to fail, and errors about symbol incompatibility will appear in the dmesg log.
	Keywords: openibd, symbol, rdma-core, driver load
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.2-2.2.0.0
	Description : When disabling KEEP_ETH_LINK_UP_P1/2 configuration while in SwitchDev mode, packets will no longer be received.
2397281	Keywords: ASAP ² ;KEEP_ETH_LINK_UP_P, SwitchDev
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.2-2.2.0.0
	Description: When an error occurs during adding a new flow rule (e.g. FW command failure in DMFS, unsupported sequence in SMFS, etc.), the driver might fail to release some resources.
2359599	Keywords: Steering, SMFS, DMFS
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.2-2.2.0.0
	Description: On EulerOS 2.0 SP9 Aarch64 with errata kernel kernel- 4.19.90-vhulk2009.2.0.h269.eulerosv2r9.aarch64, MLNX_OFED installs correctly, but driver restart fails.
2407076	Keywords: OpenEuler, Aarch64, installation,add-kernel-support
	Discovered in Release: 5.2-1.0.4.0
	Fixed in Release: 5.2-2.2.0.0
2244416	Description: Configuring "other" channels over one representor is not supported and may cause a call trace.
	Keywords: ASAP, SwitchDev, ethtool, representor
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
2209987	Description: aRFS feature (activated using "ethtool ntuple on") is disabled for kernel 4.1 or below.

Internal Reference Number	Description
	Keywords: aRFS
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the issue where clearing min_rate on all SR-IOV legacy VFs after setting min_rate to at least one of the VFs did not disable QoS min_rate.
2312015	Keywords: SR-IOV, Legacy, QoS, VF, min_rate
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Allowed installation of KMP (kmod) on Lustre kernels usingadd-kernel-supportkmp on RHEL 7.x and above systems.
2322044	Keywords: Installation, KMP, kmod, RHEL, Lustre, add-kernel-support
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: An empty /usr/src/mlnx-ofa_kernel/default is now no longer created.
2325883	Keywords: mlnx-ofa_kernel
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.2-1.0.4.0
	Description: Enabled installation of mlnx-ofa_kernel package to /bin/python3 instead of /usr/bin/python3 on RHEL 8.x systems.
2326961	Keywords: /bin/python3, mlnx-ofa_kernel, RHEL, installation
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.2-1.0.4.0
2329152	Description: Fixed the issue where installing packages through a repository failed after generation of metapackages usingadd-kernel-support. The failure occurred due to excessive and incorrect Obsolete headers in those metapackages.
	Keywords:add-kernel-support, Installation, metapackage, repository

Internal Reference Number	Description
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the issue where all IPoIB offload packets were wrongly counted as rx_csum_complete. These packets are now identified as rx_csum_unnecessary packets.
2152217	Keywords: IPoIB, rx_csum_complete, rx_csum_unnecessary
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue where mlx5dv_dr_rule_destroy segmentation fault could happen in rare cases with multiple rules on the same matcher with a different number of actions. This could happen after reusing an already deleted rule memory with less actions.
2392338	Keywords: mlx5dv_dr, software steering, RMDA-Core
	Discovered in Release: 5.1-2.3.7.1
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the OFED udev script to not modify non-NVIDIA NIC names.
2297535	Keywords: udev, SwitchDev
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue of moving to SwitchDev mode after configuring DSCP.
2333242	Keywords: DSCP, SwitchDev
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
2368330	Description: Fixed the issue where global traffic class configuration did not take effect on DC QPs.
	Keywords: DC, QP, TC
	Discovered in Release: 5.1-0.6.6.0

Internal Reference Number	Description
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue where ib_send_bw traffic frequently dropped to zero when RDMA CM was used, because of incorrect min_rnr_timer setting on the responder side. The min_rnr_timer setting is now aligned with the setting in non-RDMA CM cases.
2241983	Keywords: RDMA CM, ib_send_bw
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue where traffic did not pass over VFs with VST QinQ feature is enabled.
2334009	Keywords: VST, QinQ
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue of when traffic was sent over Geneve VLAN with Tx VLAN offload enabled and TSO or Tx csum enabled, traffic could be dropped and not sent to the wire.
2336091	Keywords: Geneve, VLAN, Tx, TSO, offload
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue that caused the firmware to restart upon installing mlnx-ofed-dpdk-upstream-libs package manually.
2317257	Keywords: Firmware, mlnx-ofed-dpdk-upstream-libs, installation
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
2365123	Description: Fixed bad utility paths in rdma-core "dracut" hooks on SLES systems, which used to result in the following errors when running "dracut" with the "add rdma" option. dracut-install: ERROR: installing '/usr/libexec/mlx4-setup.sh' dracut: /usr/lib/dracut/dracut-install -D /var/tmp/dracut.UdCOSJ/initramfs /usr/libexec/mlx4-setup.sh dracut-install: ERROR: installing '/usr/libexec/rdma-set-sriov-vf'

Internal Reference Number	Description
	dracut: /usr/lib/dracut/dracut-install -D /var/tmp/dracut.UdCOSJ/initramfs /usr/libexec/rdma-set-sriov-vf
	Keywords: rdma-core dracut
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
2379620	 Description: Fixed an issue where MLNX_OFED installation using 'yum' tool failed with the following errors. Requires: rdma-core(x86-64) = 50mlnx1-1.50218 Removing: rdma-core-22.4-1.el7.x86_64 (@anaconda) rdma-core(x86-64) = 22.4-1.el7 Obsoleted By: mlnx-ofed-all-5.0-2.1.8.0.rhel7.8.noarch (MLNX_LOCAL_REPO) Not found Updated By: rdma-core-50mlnx1-1.50218.x86_64 (MLNX_LOCAL_REPO) rdma-core(x86-64) = 50mlnx1-1.50218 Available: rdma-core-22.4-5.el7.x86_64 (base) rdma-core(x86-64) = 22.4-5.el7
	Keywords: YUM, rdma-core conflict
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the issue where traffic class value was not updated in DCT when set via sysfs.
2299982	Keywords: DCT, sysfs, TC
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed a doorbell loss issue on AMD platforms with Secure Memory Encryption (SME).
2335165	Keywords: AMD, SME
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
2355878	Description: Fixed an issue with registering memory using mlx5dv_devx_umem_reg while forking. Without this fix, applications which use fork() or similar syscalls while using a memory registered

Internal Reference Number	Description
	with umem_reg could hang due to incorrect physical page mapping. This fix requires setting the IBV_FORK_SAFE environment variable.
	Keywords: mlx5dv_dr, SW steering, RDMA core
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the issue where unbinding the device resulted in the following message being printed to the dmesg: "failed to disable DC tracer"
2288599	Keywords: Unbind
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description : Fixed the issue where the content of file /sys/class/net//statistics/multicast may have been out of date and may have displayed values lower than the real values.
2083942	Keywords: Multicast counters
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed the issue where ERSPAN protocol was available only when turning off Tx checksum offload.
2282316	Keywords: ERSPAN, TX checksum offload
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed a udev script issue which caused non-NVIDIA devices to be renamed.
2310695	Keywords: udev, naming
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.2-1.0.4.0
2334518	Description: Fixed missing representor statistics when using ifconfig.
	Keywords: SwitchDev, representor, statistics, ifconfig
	Discovered in Release: 5.1-2.5.8.0

Internal Reference Number	Description
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed wrong value of skb mark of received packets on representors.
2342348	Keywords: SwitchDev, skb mark
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed an issue which caused second port representors to be named as first port representors.
2363982	Keywords: SwitchDev, udev, representor
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.2-1.0.4.0
	Description: Fixed a kernel panic scenario that may have taken place when using sysfs to cancel the probing of VFs and performing reboot while the VFs are still managed by the mlx5 driver.
2292762	Keywords: Proved VFs
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Fixed a GPUDirect locking bug that may have caused instability and communication loss in Peer-direct applications.
2302010	Keywords: Peer-direct, GPUDirect
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.1-2.5.8.0
2298308	Description: Fixed the issue where pinned pages were not handled properly by peer flow, which resulted in ENOMEM error.
	Keywords: Peer flow
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.1-2.5.8.0
2323446	Description: Fixed the issue where num_free_callbacks counter was not functional, and querying the counters returned the value of 0 all

Internal Reference Number	Description
	the time.
	Keywords: GPUDirect
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Fixed a KMP support issue in EulerOS 2.0 SP9 with non- default kernel.
2288779	Keywords: KMP, EulerOS, OS
	Discovered in Release: 5.1-2.5.8.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Fixed the following VST issues.
2298285	 Ingress VST traffic tagging issue caused by incorrect configuration settings over VF interface. Enabled proper reset of ingress and egress legacy ACL.
	Keywords: VST
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Added command interface resilience by manually polling the async commands EQE in case of command timeout.
2247363	Keywords: Command interface, mlx5_core
	Discovered in Release: 4.2-1.2.0.0
	Fixed in Release: 5.1-2.5.8.0
2244729	Description: Enabled on-demand device memory sync to free cached memory.
	Keywords: Device memory, cache
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
2229683	Description: Eliminated long mlx5 recovery-flow delay on a VF driver when PCI interface goes down.

Internal Reference Number	Description
	Keywords: mlx5, VF, PCI
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: ibv_get_device_list returns only accessible devices now.
	Keywords: ibv_get_device_list
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: The install script will now also install the package rpm- build (that includes basic support for building RPM packages) on SLES15.x systems when trying to rebuild packages, as is the case with the rest of the distributions.
2257363	Keywords: RPM, SLES, OS
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: AF_XDP is now functional.
0044000	Keywords: AF_XDP
2244336	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
2192791	Description: Fixed the issue where packages neohost-backend and neohost-sdk were not properly removed by the uninstallation procedure and may have required manual removal before re-installing or upgrading the MLNX_OFED driver.
	Keywords: NEO-Host, SDK
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
2247404	Description: srptools installation no longer fails in case the srp_daemon service fails to start on Debian and Ubuntu systems.
	Keywords: SRP, srptool, srp_daemon, Debian, Ubuntu, OS
	Discovered in Release: 5.1-0.6.6.0

Internal Reference Number	Description
	Fixed in Release: 5.1-2.5.8.0
	Description: Fixed the issue where during MLNX_OFED installation, warning messages related to modules iw_cxgb3 and iw_nes may have appeared in the log.
2248201	Keywords: SLES, RHEL, KMP, weak updates, kmod
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
2201475	Description : Simplified the sanity test that verifies which systems must use python3 and not python2, so that when rebuilding the kernel packages usingadd-kernel-support, the systems based on CentOS 8 are detected.
2261475	Keywords: CentOS 8, RHEL 8, OS, sanity test, python
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Simplified sanity tests for building MLNX_OFED/MLNX_EN with kernel v5.7.11 or newer.
2288099/2257	Keywords: Kernel, sanity test, installation
445	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: MLNX_OFED/MLNX_EN installer will no longer automatically remove the spdk package when installing the driver.
2288783	Keywords: spdk, installation
	Discovered in Release: 5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Marked additional packages as obsolete so they would be removed when installing the HPE-roce mlnx-ofa-kernel package.
2292273	Keywords: RPM, HPE-roce
	Discovered in Release: 5.1-1.0.7.1
	Fixed in Release: 5.1-2.5.8.0

Internal Reference Number	Description
	Description: Fixed VM driver to avoid command timeouts when the Hypervisor disables the VF's PCI interface. This reduces the device removal time to 2 seconds.
2265094	Keywords: VF disable
	Discovered in Release:5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Added back rx_dct_connect counter that queries the number of received connection requests for the associated DCTs.
2272539	Keywords: rx_dct_connect, DCT
	Discovered in Release:5.1-0.6.6.0
	Fixed in Release: 5.1-2.5.8.0
	Description: Get/Set Forward Error Correction FEC configuration is not supported on ConnectX-6 HCAs with 200Gbps speed rate.
1731939	Keywords: Forward Error Correction, FEC, 200Gbps
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 5.1-0.6.6.0
	Description: Setting VF VLAN, state and spoofchk using ip link tool is not supported in SwitchDev mode.
1980884	Keywords: ASAP, ip tool, VF, SwitchDev
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.0-1.0.0.0
2117845	Description: Relaxed ordering memory regions are not supported when working with CAPI. Registering memory region with relaxed ordering while CAPI enabled will result in a registration failure.
	Keywords: Relaxed ordering, memory region, MR, CAPI
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
2118956	Description: mlx5dv_dr API does not support sub functions (SFs) as destination actions.

Internal Reference Number	Description
	Keywords: mlx5dv_dr, sub functions, SF
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Userspace Software Steering using mlx5dv_dr API support on ConnectX-6 Dx adapter cards is now at GA level.
2097045	Keywords: Software Steering, SW, mlx5dv_dr, ConnectX-6 Dx
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed a sporadic reporting bandwidth issue in case of running withrun_infinitely flag.
2132332	Keywords: perftest, bandwidth
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Optimized XRC target lookup by modifying the locking scheme to enable multiple readers and changing the linked list that holds the QPs to xarray.
2151658	Keywords: XRC, QP, xarray
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed a driver issue that led to panic after DPDK application crashes.
2196118	Keywords: DPDK, panic
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed an issue of a crash when attempting to access roce_enable sysfs in unprobed VFs.
2245228	Keywords: roce_enable, unprobed VFs
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0

Internal Reference Number	Description
2061294	Description: Fixed a race of commands executed by command interface in parallel to AER recovery causing the kernel to crash.
	Keywords: mlx5e, AER
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed an issue in MLNX_OFED build system that broke RPM sign process for random packages, all RPMs are now signed properly.
2131951	Keywords: RPM, sign
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.1-0.6.6.0
	Description: If Openibd was configured to enable the SRP daemon, it now also enables srp_daemon from rdma-core.
2143067	Keywords: Openibd, SRP daemon, srp_daemon, rdma-core
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Regenerated package repository in the correct location after rebuilding the kernel using add-kernel-support. This allows for installing the newly generated packages with a package manager.
2143094	Keywords: add-kernel-support, RPM, deb
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.1-0.6.6.0
2172130	Description: Fixed an issue with metadata packages generation in the eth-only directory. This allows using the directory as a repository for package managers.
	Keywords: Metadata packages
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
2214543	Description: Moved ibdev2netdev script from /usr/bin to /usr/sbin in the RPM package to avoid package conflict with RHEL 8 and

Internal Reference Number	Description
	consequent MLNX_OFED installation failure on some systems.
	Keywords: ibdev2netdev, RPM, RHEL, RedHat
	Discovered in Release:
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed an issue where Rx port buffers cell size was wrong, leading to wrong buffers size reported by mlnx_qos/netdev qos/buffer_size sysfs.
2211311	Keywords: mlx5e, RX buffers, mlnx_qos
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where ethtoolshow-fec/get-fec were not supported over ConnectX-6 and ConnectX-6 Dx adapter cards.
2111349	Keywords: Ethtool, ConnectX-6 Dx
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed an issue related to mlx5 command interface that in some scenarios caused the driver to hang.
2165668	Keywords: ConnectX-5, mlx5, panic
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where IPsec crypto offloads did not work when ESN was enabled.
2119984	Keywords: IPsec, ESN
	Discovered in Release: 5.0-2.1.8.0
	Fixed in Release: 5.1-0.6.6.0
1630228	Description: Fixed the issue where tunnel stateless offloads were wrongly forbidden for E-Switch manager function.
	Keywords: Stateless offloads cap
	Discovered in Release: 4.6-1.0.1.1

Internal Reference Number	Description
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where dump flows were not supported and may have been corrupted when using tc tool with connection tracking rules.
2089996	Keywords: ASAP, iproute2, tc, connection tracking
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue of when one of the LAG slaves went down, LAG deactivation failed, ultimately causing bandwidth degradation.
2094216	Keywords: RoCE LAG
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
2133778	Description: The mlx5 driver maintains a subdirectory for every open eth port in /sys/kernel/debug/. For the default network namespace, the sub-directory name is the name of the interface, like "eth8". The new convention for the network interfaces moved to the non-default network namespaces is the interfaces name followed by "@" and the port's PCI ID. For example: "eth8@0000:af:00.3".
	Keywords: Namespace
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where in RPM-based OSs with non-default kernels, using repositories after re-creating the installer (usingadd-kernel-support) would result in improper installation of the drivers.
2076546	Keywords: Installation, OS
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.1-0.6.6.0
2114957	Description: Fixed the issue where MLNX_OFED installation may have depended on python2 package even when attempting to install it on OSs whose default package is python3.
	Keywords: Installation, python

Internal Reference Number	Description
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
2122684	Description: Fixed the issue where OFED uninstallation resulted in the removal of dependency packages, such as qemu-system-* (qemu-system-x86).
	Keywords: Uninstallation, dependency, qemu-system-x86
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Added KMP ability to install MLNX_OFED Kernel modules on SLES12 SP5 and SLES15 kernel maintenance updates.
2135476	Keywords: KMP, SLES, kernel
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
2143258	Description: Fixed a typo in perftest package where help messages wrongly displayed the conversion result between Gb/s and MB/s (20^2 instead of 2^20).
	Keywords: perftest
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
2149577	Description: Fixed the issue where openibd script load used to fail when esp6_offload module did not load successfully.
	Keywords: openibd, esp6_offload
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
2163879	Description: Added dependency of package mpi-selectors on perl- Getopt-Long system package. On minimal installs of RPM-based OSs, installing mpi-selectors will also install the required system package perl-Getopt-Long.
	Keywords: Dependency, perl-Getopt-Long
	Discovered in Release: 4.7-3.2.9.0

Internal Reference Number	Description
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where injecting EEH may cause extra Kernel prints, such as: "EEH: Might be infinite loop in mlx5_core driver".
2119017	Keywords: EEH, kernel
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0
	Description: Fixed the issue where in certain rare scenarios, due to Rx page not being replenished, the same page fragment mistakenly became assigned to two different Rx descriptors.
2107532	Keywords: Memory corruption, Rx page recycle
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.0-2.1.8.0
	Description: Fixed the issue where ibsim was missing after OFED installation.
2116234	Keywords: ibsim, installation
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
	Description: Fixed an issue where ucx-kmem was missing after OFED installation.
2116233	Keywords: ucx-kmem, installation
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
	Description: Fixed a dependency issue between systemd and RDMA-Core.
2109716	Keywords: Dependency, RDMA-Core
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2107776	Description: Fixed a driver load issue with Errata-kernel on SLES15 SP1.

Internal Reference Number	Description
	Keywords: Load, SLES, Errata
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2105536	Description: Fixed an issue in the Hairpin feature which prevented adding hairpin flows using TC tool.
	Keywords: Hairpin, TC
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
	Description: Fixed the issue where WQ queue flushing was not handled properly in the event of EEH.
2090321	Keywords: WQ, EEH
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.0-2.1.8.0
	Description: Fixed a rare kernel crash scenario when exiting an application that uses RMPP mads intensively.
2076311	Keywords: MAD RMPP
	Discovered in Release: 4.0-1.0.1.0
	Fixed in Release: 5.0-2.1.8.0
2094545	Description: Fixed the issue where perftest applications (ib_read_*, ib_write_* and others) supplied with MLNX_OFED v5.0 and above did not work correctly if corresponding applications on another side of client-server communication were supplied with previous versions of MLNX_OFED due to an interoperability issue.
	Keywords: perftest, interoperability
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2096998	Description: Fixed the issue where NEO-Host could not be installed from the MLNX_OFED package when working on Ubuntu and Debian OSs.
	Keywords: NEO-Host, Ubuntu, Debian

Internal Reference Number	Description
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2094012	Description: Fixed the issue where MLNX_OFED installation failed to upgrade firmware version on ConnectX-6 Dx NICs with secure-fw.
	Keywords: ConnectX-6 Dx, installation, firmware, NIC
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2057076	Description: Added support for installing MLNX_OFED usingadd-kernel- support option over RHEL 8 OSs.
	Keywords:add-kernel-support, installation, RHEL
	Discovered in Release: 5.0-1.0.0.0
	Fixed in Release: 5.0-2.1.8.0
2090186	Description: Fixed a possible kernel crash scenario when AER/slot reset in done in parallel to user space commands execution.
	Keywords: mlx5_core, AER, slot reset
	Discovered in Release: 4.3-1.0.1.0
	Fixed in Release: 5.0-2.1.8.0
	Description: Added missing ECN configuration under sysfs for PFs in SwitchDev mode.
2093410	Keywords: sysfs, ASAP, SwitchDev, ECN
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.0-2.1.8.0
1731005	Description: Fixed the issue where MLNX_OFED v4.6 YUM and Zypper installations failed on RHEL8.0, SLES15.0 and PPCLE OSs.
	Keywords: YUM, Zypper, installation, RHEL, RedHat, SLES, PPCLE
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 5.0-1.0.0.0
1779150	Description: Fixed the issue of when upgrading the MLNX_OFED version over SLES 15 SP0 and SP1 OSs on PPCLE platforms, it might

Internal Reference Number	Description
	have failed due to an isert-kmp-default issue.
	Keywords: Installation, SLES, PPCLE
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 5.0-1.0.0.0
	Description: Fixed the issue of when using the RDMA statistics feature and attempting to unbind a QP from a counter, not including the counter-id as an argument in the CLI would have resulted in a segmentation fault.
1897199	Keywords: RDMA, QP, segfault, unbinding
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 5.0-1.0.0.0
	Description: Fixed the issue of when firmware response time to commands became very long, some commands failed upon timeout. The driver may have then triggered a timeout completion on the wrong entry, leading to a NULL pointer call trace.
1916029	Keywords: Firmware, timeout, NULL
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.0-1.0.0.0
	Description: Added driver support for kernels with the old XDP_REDIRECT infrastructure that uses the following NetDev operations: .ndo_xdp_flush and .ndo_xdp_xmit.
2036394	Keywords: XDP_REDIRECT, Soft lockup
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.0-1.0.0.0
1973238	Description: Fixed the issue where ib_core unload may fail on Ubuntu 18.04.2 OS with the following error message: "Module ib_core is in use"
	Keywords: ib_core, Ubuntu, ibacm
	Discovered in Release: 4.7-3.2.9.0
	Fixed in Release: 5.1-0.6.6.0

Internal Reference Number	Description	
	Description: Fixed an issue where the usage ofexcludedocs Open MPI RPM option resulted in the removal of non-documentation related files.	
2072871	Keywords:excludedocs, Open MPI, RPM	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Legacy mlnx-libs are now installed by default on SLES11 SP3 OS, as building MLNX_OFED on RDMA-Core based packages with this OS is not supported.	
2060216	Keywords: mlnx-libs, SLES, RDMA-Core	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
2072884	Description: Removed all cases of automated loading of MLNX_OFED kernel modules outside of openibd to preserve the startup process of previous MLNX_OFED versions. These loads conflict with openibd, which has its own logic to overcome issues. Such issues can be inbox driver load instead of MLNX_OFED, or module load with wrong parameter value. They might also load modules while openibd is trying to unload the driver stack.	
	Keywords: Installation, openibd, RDMA-Core	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Disabled automated loading of some modules through udev triggers to preserve the startup process of previous MLNX_OFED versions.	
2052037	Keywords: Installation, udev, RDMA-Core	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
2022634	Description: Fixed a typo in the packages build command line which could cause the installation of MLNX_OFED on SLES OSs to fail when using the optionwithout-depcheck.	
	Keywords: Installation, SLES	

Internal Reference Number	Description	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed the issue where uninstallation of MLNX_OFED would hang due to a bug in the package dependency check.	
2022619	Keywords: Uninstallation, dependency	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: ibdump is now provided with the default rdma-core-based build.	
1995843	Keywords: ibdump, RDMA-Core	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Proper package dependencies are now set on Debian and Ubuntu libibverbs-dev package that is generated from RDMA-Core.	
1995631	Keywords: Dependency, libibverbs, RDMA-Core	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Reference count (refcount) for RDMA connection ID (cm_id) was not incremented in rdma_resolve_addr() function, resulting in a cm_id use-after-free access. A fix was applied to increment the cm_id refcount.	
2047221	Keywords: rdma_resolve_addr(), cm_id	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed a race condition which caused kernel panic when moving two ports to SwitchDev mode at the same time.	
2045181	Keywords: ASAP, SwitchDev, race	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	

Internal Reference Number	Description	
	Description: Allowed accessing sysfs hardware counters in SwitchDev mode.	
2004488	Keywords: ASAP, hardware counters, sysfs, SwitchDev	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	
2030943	Description: Function smp_processor_id() is called in the RX page recycle flow to determine the core to run on. This is intended to run in NAPI context. However, due to a bug in backporting, the RX page recycle was mistakenly called also in the RQ close flow when not needed.	
	Keywords: Rx page recycle, smp_processor_id	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed an issue where port link state was automatically changed (without admin state involvement) to "UP" after reboot.	
2074487	Keywords: Link state, UP	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed an issue where RDMA CM connection failed when port space was small.	
2064711	Keywords: RDMA CM	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Traffic mirroring with OVS offload and non-offload over VxLAN interface is now supported. Note : For kernel 4.9, make sure to use a dedicated OVS version.	
2076424	Keywords: VxLAN, OVS	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	

Internal Reference Number	Description	
	Description: Fixed the issue of when working with VF LAG while the bond device is in active-active mode, running fwreset would result in unequal traffic on both PFs, and PFs would not reach line rate.	
1828321	Keywords: VF LAG, bonding, PF	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Installing OFED withwith-openvswitch flag no longer requires manual removal of the existing Open vSwitch.	
1975293	Keywords: OVS, Open vSwitch, openvswitch	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
1939719	Description: Fixed an issue of when running openibd restart after the installation of MLNX_OFED on SLES12 SP5 and SLES15 SP1 OSs with the latest Kernel (v4.12.14) resulted in an error that the modules did not belong to that Kernel. This was due to the fact that the module installed by MLNX_OFED was incompatible with new Kernel's module.	
	Keywords: SLES, operating system, OS, installation, Kernel, module	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
0001000	Description: Fixed an issue of when bond was created over VF netdevices in SwitchDev mode, the VF netdevice would be treated as representor netdevice. This caused the mlx5_core driver to crash in case it received netdevice events related to bond device.	
2001966	Keywords: PF, VF, SwitchDev, netdevice, bonding	
	Discovered in Release: 4.7-3.2.9.0	
	Fixed in Release: 5.0-1.0.0.0	
1816629	Description: Fixed an issue where following a bad affinity occurrence in VF LAG mode, traffic was sent after the port went up/down in the switch.	
	Keywords: Traffic, VF LAG	
	Discovered in Release: 4.6-1.0.1.1	

Internal Reference Number	Description	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Added support for VLAN header rewrite on CentOS 7.2 OS.	
1718531	Keywords: VLAN, ASAP, switchdev, CentOS 7.2	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed the issue where adding VxLAN decapsulation rule with enc_tos and enc_ttl failed.	
1556337	Keywords: VxLAN, decapsulation	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 5.0-1.0.0.0	
	Description: Fixed the issue where MLNX_OFED installation over SLES15 SP1 ARM OSs failed unlessadd-kernel-support flag was added to the installation command.	
1921799	Keywords: SLES, installation	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 4.7-3.2.9.0	
	Description: Fixed a race condition that resulted in kernel panic when running IPoIB traffic in Connected mode.	
1949260	Keywords: IPoIB	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 4.7-3.2.9.0	
	Description: Fixed wrong EEPROM length for small form factor (SFF) 8472 from 256 to 512 bytes.	
1973828	Keywords: EEPROM, SFF	
	Discovered in Release: 4.7-1.0.0.1	
	Fixed in Release: 4.7-3.2.9.0	
1915553	Description: Fixed the issue where errno field was not sent in all error flows of ibv_reg_mr API.	

Internal Reference Number	Description
	Keywords: ibv_reg_mr
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 4.7-3.2.9.0
	Description: Fixed the issue where mlx5 IRQ name did not change to express the state of the interface.
1970901	Keywords: Ethernet, PCIe, IRQ
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 4.7-3.2.9.0
	Description: Udaddy application is now functional in Legacy mode.
1015507	Keywords: Udaddy, MLNX_OFED legacy, RDMA-CM
1915587	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 4.7-3.2.9.0
	Description: Added support for E-Switch (SR-IOV Legacy) mode in RHEL 7.7 OSs.
1931421	Keywords: E-Switch, SR-IOV, RHEL, RedHat
	Discovered in Release: 4.7-1.0.0.1
	Fixed in Release: 4.7-3.2.9.0
	Description: Fixed the issue of when XDP_REDIRECT fails, pages got double-freed due to a bug in the refcnt_bias feature.
1945411/1839	Keywords: XDP, XDP_REDIRECT, refcnt_bias
353	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 4.7-3.2.9.0
	Description: Fixed the issue where NVIDIA Firmware Tools (MFT) package was missing from Ubuntu v18.04.2 OS.
1715789	Keywords: MFT, Ubuntu, operating system
	Discovered in Release: 4.6-1.0.1.1
	Fixed in Release: 4.7-1.0.0.1

Internal Reference Number	Description	
	Description : Fixed an issue where IPoIB Tx queue may get stuck, leading to timeout warnings in dmesg.	
1547200	Keywords: IPoIB	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 4.7-1.0.0.1	
	Description: Fixed the issue of when disabling one port on the Server side, VF-LAG Tx Affinity would not work on the Client side.	
1817636	Keywords: VF-LAG, Tx Affinity	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 4.7-1.0.0.1	
	Description: When configuring the Time-stamping feature, CQE compression will be disabled. This fix entails the removal of a warning message that appeared upon attempting to disable CQE compression when it has already been disabled.	
1800525	Keywords: Time-stamping, CQE compression	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 4.7-1.0.0.1	
	Description : Fixed the issue where software reset may have resulted in an order inversion of interface names.	
1431282	Keywords: Software reset	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.7-1.0.0.1	
	Description: Server reboot may result in a system crash.	
1042020	Keywords: reboot, crash	
1843020	Discovered in Release: 4.2-1.2.0.0	
	Fixed in Release: 4.7-1.0.0.1	
1734102	Description: Fixed the issue where Ubuntu v16.04.05 and v16.04.05 OSs could not be used with their native kernels.	
	Keywords: Ubuntu, Kernel, OS	

Internal Reference Number	Description	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 4.7-1.0.0.1	
	Description: VF mirroring offload is now supported.	
1011072	Keywords: ASAP ² , VF mirroring	
1811973	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 4.7-1.0.0.1	
	Description: The number of guaranteed counters per VF is now calculated based on the number of ports mapped to that VF. This allows more VFs to have counters allocated.	
1841634	Keywords: Counters, VF	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.7-1.0.0.1	
	Description: Installing MLNX_OFED on RHEL 7.6 OSs platform x86_64 and RHEL 7.6 ALT OSs platform PPCLE using YUM is now supported.	
1758983	Keywords: RHEL, RedHat, YUM, OS, operating system	
	Discovered in Release: 4.6-1.0.1.1	
	Fixed in Release: 4.7-1.0.0.1	
	Description: Fixed the issue where RDMA connection persisted even after dropping the network interface.	
1523548	Keywords: Network interface, RDMA	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.6-1.0.1.1	
1712870	Description: Fixed the issue where small packets with non-zero padding were wrongly reported as "checksum complete" even though the padding was not covered by the csum calculation. These packets now report "checksum unnecessary". In addition, an ethtool private flag has been introduced to control the "checksum complete" feature: ethtoolset-priv-flags eth1 rx_no_csum_complete on/off	
	Keywords: csum error, checksum, mlx5_core	

Internal Reference Number	Description	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 4.6-1.0.1.1	
	Description : Fixed the wrong wording in the FW tracer ownership startup message (from "FW Tracer Owner" to "FWTracer: Ownership granted and active").	
1648597	Keywords: FW Tracer	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 4.6-1.0.1.1	
	Description : Fixed the issue where GID entries referenced to by a certain user application could not be deleted while that user application was running.	
1581631	Keywords: RoCE, GID	
	Discovered in Release: 4.5-1.0.1.0	
	Fixed in Release: 4.6-1.0.1.1	
	Description : Fixed the issue where MLNX_OFED could not be installed on RHEL 7.x Alt OSs using YUM repository.	
1368390	Keywords: Installation, YUM, RHEL	
	Discovered in Release: 4.3-3.0.2.1	
	Fixed in Release: 4.6-1.0.1.1	
	Description : Fixed an issue of when the number of channels configured was less than the number of CPUs available, part of the CPUs would not be used by Tx queues.	
1531817	Keywords: Performance, Tx, CPU	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.5-1.0.1.0	
1571977	Description : Fixed an issue of when the same CQ is connected to some QPs with SRQ and some without, wrong <i>wr_id</i> might be reported by <i>ibv_poll_cq</i> .	
	Keywords : libmlx5, wr_id	
	Discovered in Release: 4.4-1.0.0.0	

Internal Reference Number	Description	
	Fixed in Release: 4.5-1.0.1.0	
	Description : Fixed the issue where IB port link used to flap due to MAD heartbeat response delay when using new CQ API.	
1380135	Keywords: IB port link, CQ API, MAD heartbeat	
	Discovered in Release: 4.2-1.2.0.0	
	Fixed in Release: 4.5-1.0.1.0	
	Description : Fixed the issue where establishing TCP connection took too long due to failure of SA PathRecord query callback handler.	
1498931	Keywords: TCP, SA PathRecord	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.5-1.0.1.0	
	Description : Fixed the issue where lack of high order allocations caused driver load failure. All high order allocations are now changed to order-O allocations.	
1514096	Keywords: mlx5, high order allocation	
	Discovered in Release: 4.0-2.0.2.0	
	Fixed in Release: 4.5-1.0.1.0	
	Description : Fixed a backport issue on some OSs, such as RHEL v7.x, where mlx5 driver would support <i>ip link set DEVICE vf NUM rate</i> <i>TXRATE</i> old command, instead of <i>ip link set DEVICE vf NUM</i> <i>max_tx_rate TXRATE min_tx_rate TXRATE</i> new command.	
1524932	Keywords: mlx5 driver	
	Discovered in Release: 4.0-2.0.2.0	
	Fixed in Release: 4.5-1.0.1.0	
	Description : Fixed the issue of when performing configuration changes, mlx5e counters values were reset.	
1498585	Keywords: Ethernet counters	
	Discovered in Release: 4.0-2.0.2.0	
	Fixed in Release: 4.5-1.0.1.0	

Internal Reference Number	Description	
1425027	Description : Fixed the issue where attempting to establish a RoCE connection on the default GID or on IPv6 link-local address might have failed when two or more netdevices that belong to HCA ports were slaves under a bonding master. This might also have resulted in the following error message in the kernel log: "ib_cache_gid_add: unable to add gid fe80:0000:0000:0000:f652:14ff:fe46:7391 error=-28".	
	Keywords: RoCE, bonding	
	Discovered in Release: 4.4-1.0.0.0	
	Fixed in Release: 4.5-1.0.1.0	

User Manual Revision History

Releas e	Date	Description
24.07	August 2024	No changes to the User Manual.
24.04	April 2024	Added <u>SR-IOV Live Migration</u>
23.10	 Added IPsec Full Offload Added a note to NVME-oF - NVM Express over Fabrication 	
23.07	September 2023	No changes to the User Manual.
	August 2023	 Updated "Interrupt Request (IRQ) Naming" under <u>Ethernet Interface</u>. Removed the "IPsec Crypto Offload" section from the <u>Ethernet Network</u> chapter. Removed all "Connected Mode" indications, and the "IPoIB Mode Setting" section from <u>IP over InfiniBand (IPoIB)</u>. Removed Running OpenSM, and updated Running OpenSM as Daemon under <u>NVIDIA SM</u>.
23.04	May 2023	 Removed "Installing NEO-Host Using mInxofedinstall Script" subsection from the <u>Installing the Driver</u> section

Releas e	Date	Description
		 Added note under "<u>Enrolling NVIDIA's x.509 Public Key</u> <u>On your Systems</u>" subsection in the <u>UEFI Secure Boot</u> section Updated the <u>NVIDIA SM</u> (OpenSM) section
	Jun 2023	 Added an important note to <u>IPsec Crypto Offload</u>
5.7	August 2022	 Added <u>Out of Order (OOO)</u> under RoCE section Added section describing <u>Dumping Steering Info</u>
5.6	April 2022	 Added <u>OpenFlow Meters</u> section Added <u>Installation on Community Operating Systems</u> section

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	24.07-0.6.1.0	 <u>License</u> <u>3rd Party Notice</u>
Firmware	xx.42.1000	 HCA Firmware EULA 3rd Party Unify Notice License
MFT	4.29.0-131	 <u>3rd Party Notice</u> <u>License</u>
Clusterkit	1.14	 <u>License</u> <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
DPCP	1.1.48	 <u>License</u> <u>3rd Party Notice</u>
VMA	9.8.60	 <u>License</u> <u>3rd Party Notice</u>
XLIO	3.31.2	 <u>License</u> <u>3rd Party Notice</u>
HCOLL	4.8.3228	 <u>License</u> <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
SHARP	3.8.0	• <u>License</u>

Product	Version	Legal Notices and 3rd Party Licenses
		 <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
ibutils2	2.18	 <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
OpenSM	5.20.0	 <u>3rd Party Unify Notice</u> <u>3rd Party Notice</u>
mpitests	3.2.24	 <u>License</u> <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
UCX	1.17	 <u>License</u> <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>
ОМРІ	4.1	 <u>License</u> <u>3rd Party Notice</u> <u>3rd Party Unify Notice</u>

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF

ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2024, NVIDIA. PDF Generated on 10/22/2024