



Connect. Accelerate. Outperform.™

Mellanox WinOF VPI Release Notes

Rev 4.80.50000

Last Updated: November 23, 2014

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
 350 Oakmead Parkway Suite 100
 Sunnyvale, CA 94085
 U.S.A.
www.mellanox.com
 Tel: (408) 970-3400
 Fax: (408) 970-3403

Mellanox Technologies, Ltd.
 Beit Mellanox
 PO Box 586 Yokneam 20692
 Israel
www.mellanox.com
 Tel: +972 (0)74 723 7200
 Fax: +972 (0)4 959 3245

© Copyright 2014. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, MetroX®, MLNX-OS®, TestX®, PhyX®, ScalableHPC®, SwitchX®, UFM®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

ExtendX™, FabricIT™, HPC-X™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Virtual Modular Switch™, MetroDX™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
List of Tables	5
Chapter 1 Overview	6
1.1 Mellanox WinOF Rev 4.80.50000 New Features	6
1.2 WinOF VPI Package Contents	6
1.3 Supported Native Operating System Versions	7
1.4 Supported Network Adapter Cards	8
1.4.1 Firmware Versions	8
Chapter 2 Changes and Major New Features	9
2.1 New Features, Changes and Fixes in v4.80.50000 from v4.70	9
2.2 New Features, Changes and Fixes in Version 4.70	10
2.2.1 New Features, Changes and Fixes in v4.70.50050 from v4.70.50040	10
2.2.2 New Features, Changes and Fixes in v4.70.50040 from v4.70.50000	10
2.2.3 New Features, Changes and Fixes in v4.70.50000 from v4.61	11
2.3 New Features, Changes and Fixes in v4.61 From v4.60	12
2.4 New Features, Changes and Fixes in Version 4.60	13
2.4.1 New Features, Changes and Fixes in v4.60.17738 from v4.60.17718	13
2.4.2 New Features, Changes and Fixes in v4.60.17718 from v4.55	13
2.5 New Features, Changes and Fixes in v4.55 From v4.40	15
2.6 New Features, Changes and Fixes in v4.40 From v4.30	16
2.7 New Features, Changes and Fixes in v4.30 From v4.20	19
2.8 New Features, Changes and Fixes in v4.2 From v3.2.0	20
Chapter 3 Features/Support to be Deprecated in the Next Release 22	
Chapter 4 Beta Features	23
Chapter 5 Unsupported Functionality/Features	24
5.1 ConnectX®-2 Adapter Limitations	24
Chapter 6 Known Issues	25
6.1 Generic Issues	25
6.2 InfiniBand Issues	27
6.3 Ethernet Issues	29
6.4 Quality of Service Issues	31
6.5 RoCE NVGRE Issues	31
6.6 Performance Issues	32
6.7 IPoIB Issues	32
6.8 Hyper-V Issues	33
6.9 Installation / Upgrade Issues	34
6.10 Utilities Issues	37
6.11 CIM/WMI Issues	38
Chapter 7 API Changes	39

7.1	API Changes in WinOF Rev 4.80.50000.....	39
7.2	API Changes in WinOF Rev 4.70	39
7.3	API Changes in WinOF Rev 4.60	39
7.4	API Changes in WinOF Rev 4.55	40
7.5	API Changes in WinOF Rev 4.40	40

List of Tables

Table 1:	Firmware Versions	8
Table 2:	New Features, Changes and Fixes in v4.80.50000	9
Table 3:	New Features, Changes and Fixes in v4.70.50040	10
Table 4:	New Features, Changes and Fixes in v4.70.50040	10
Table 5:	New Features, Changes and Fixes in v4.70.50000	11
Table 6:	New Features, Changes and Fixes in v4.61	12
Table 7:	New Features, Changes and Fixes in v4.60.17738	13
Table 8:	New Features, Changes and Fixes in v4.60.17718	13
Table 9:	New Features, Changes and Fixes in v4.55	15
Table 10:	New Features, Changes and Fixes in v4.40	16
Table 11:	New Features, Changes and Fixes in v4.3	19
Table 12:	New Features, Changes and Fixes in v4.2	20
Table 13:	Beta Features	23
Table 14:	Generic Issues	25
Table 15:	InfiniBand Issues	27
Table 16:	Ethernet Issues	29
Table 17:	Quality of Service Issues	31
Table 18:	RoCE NVGRE Issues	31
Table 19:	Performance Issues	32
Table 20:	IPoIB Issues	32
Table 21:	Hyper-V Issues	33
Table 22:	Installation / Upgrade Issues	34
Table 23:	Utilities Issues	37
Table 24:	CIM/WMI Issues	38
Table 25:	WinOF Known Issues	38
Table 26:	API Changes in WinOF Rev 4.80	39
Table 27:	API Changes in WinOF Rev 4.70	39
Table 28:	API Changes in WinOF Rev 4.60	39
Table 29:	API Changes in WinOF Rev 4.55	40
Table 30:	API Changes in WinOF Rev 4.40	40

1 Overview

These are the release notes for Mellanox WinOF Rev 4.80.50000 VPI drivers.

The driver provides improved performance and additional functionality compared to the Inbox driver provided in Windows 2012 and Windows 2012 R2. Therefore, Mellanox strongly recommends updating to its latest driver.

1.1 Mellanox WinOF Rev 4.80.50000 New Features

Added support for the following:

- InfiniBand:
 - IPoIB multiple Pkey support (beta level)
- Ethernet/ RoCE:
 - RoCE IP Based compatible with the MLNX_OFED Linux driver RoCE mode
 - Read active RoCE configuration from hardware¹
 - RSS in UDP (enabled by default)
 - 56 Gb/s Ethernet
- SR-IOV over InfiniBand:
 - SR-IOV InfiniBand over KVM Hypervisor
 - IPv6 support of IPoIB in an SR-IOV guest OS over KVM is at beta level
 - Non-default P_Key in VM
 - Reset Flow (+SR-IOV)¹
- Generic:
 - Windows 7 Client (64 bit only)
 - Windows 8.1 Client (64 bit only)

1.2 WinOF VPI Package Contents

The Mellanox WinOF Rev 4.80.50000 for Windows package contains the following components:

- Core and ULPs:
 - IB HCA low-level drivers (mlx4)
 - IB Access Layer (IBAL)
 - Ethernet driver (ETH)
 - IP over InfiniBand (IPoIB)
 - NetworkDirect (ND)
- Mellanox LBFO driver for Windows 2008 R2

1. Requires firmware v2.32.5100

- Utilities:
 - OpenSM: InfiniBand Subnet Manager is provided as a sample code. The sample code is intended to allow users to test or bring-up the InfiniBand fabric without a management console / switch (to get started).
For cluster production environments, Mellanox's recommendation is to use a Managed Switch or the UFM-SDN Appliance.
 - Low level performance tools
- InfiniBand Diagnostics tools
- CIM, PowerShell, and WMI support
- Software Development Kit (SDK)
- Documentation

1.3 Supported Native Operating System Versions

The following describes the supported operating systems and their roles in a virtualization environment.

- **Native:**
 - Windows Server 2008 R2 (64 bit only)
 - Windows Server 2012 (64 bit only)
 - Windows Server 2012 R2 (64 bit only)
 - Windows 7 Client (64 bit only)
 - Windows 8.1 Client (64 bit only)
- **Hyper-V (non-SR-IOV):**
 - Hypervisor OS: Windows Server 2008 R2 (64 bit only)
 - Virtual machine OS: Windows Server 2008 R2 (64 bit only)
 - Hypervisor OS: Windows Server 2012 (64 bit only)
 - Virtual machine OS: Windows Server 2008 R2 (64 bit only), and Windows Server 2012 (64 bit only)
 - Hypervisor OS: Windows Server 2012 R2 (64 bit only)
 - Virtual machine OS: Windows Server 2008 R2 (64 bit only), Windows Server 2012 (64 bit only), and Windows Server 2012 R2 (64 bit only)
- **SR-IOV Ethernet:**
 - Hypervisor OS: Windows Server 2012 R2 (64 bit only)
 - Virtual machine OS: Windows Server 2012 (64 bit only) and Windows Server 2012 R2 (64 bit only)
- **SR-IOV InfiniBand:**
 - Hypervisor OS: KVM with MLNX_OFED 2.3 or later
 - Virtual machine OS: Windows Server 2008 R2 (64 bit only), Windows Server 2012 (64 bit only), and Windows Server 2012 R2 (64 bit only)

1.4 Supported Network Adapter Cards

Mellanox WinOF Rev 4.80.50000 supports the following Mellanox network adapter cards:

- ConnectX®-3 Pro and ConnectX®-3 Pro EN
 - 10, 40 and 56 Gb/s InfiniBand (IB)
 - 10, 40 and 56 Gb/s Ethernet
- ConnectX®-3 and ConnectX®-3 EN
 - 10, 40 and 56 Gb/s InfiniBand (IB)
 - 10, 40 and 56 Gb/s Ethernet
- ConnectX®-2 and ConnectX®-2 EN
 - 10 and 40 Gb/s InfiniBand (IB)
 - 10 Gb/s Ethernet

1.4.1 Firmware Versions

Mellanox WinOF Rev 4.80.50000 provides the following firmware for Mellanox NICs:

Table 1 - Firmware Versions

HCA	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3 Pro / ConnectX®-3 Pro EN	Rev 2.32.5100	Rev 2.31.5050
ConnectX®-3 / ConnectX®-3 EN	Rev 2.32.5100	Rev 2.31.5050
ConnectX®-2 / ConnectX®-2 EN	Rev 2.9.1000	Rev 2.9.1000

2 Changes and Major New Features

2.1 New Features, Changes and Fixes in v4.80.50000 from v4.70



This package version is 4.80.50000. The package contains the following versions of components.

- Bus, eth, IPoIB and mux drivers version is 4.80.10388
- The CIM provider version is 4.80.10388

Table 2 - New Features, Changes and Fixes in v4.80.50000

Category	Description
Installation/ Upgrade	<ul style="list-style-type: none"> • Added check for administrator privileges during installation • Added support for installation in silent mode without execution of <code>perf_tune</code> • Fixed installation stuck when Remote Desktop Session Host Windows Installer RDS compatibility is enabled
Generic	<ul style="list-style-type: none"> • Changed Reset Flow (+SR-IOV)-enabled only if no user space application is running and depends on the registry key: AllowResetOnError setting) • Changed the number of supported QPs in a multicast group from hard coded value to firmware capabilities dependent • Fixed driver load failure in machines with 1 TB memory and above
IPoIB	<ul style="list-style-type: none"> • Added multiple P_Key support (beta level) • Added IPoIB SR-IOV over KVM and ESX Hypervisors (for both full and partial membership) • Added support for LID change event • Added enhancements in <code>part_man</code> for the multiple Pkey support • Changed IPv6 “all dhcp servers” mcast to be persistent • Fixed rare cases of driver hang following a Subnet Manager failover event • Fixed stability issues
Ethernet	<ul style="list-style-type: none"> • Added RSS in UDP (enabled by default) • Added 56 GbE (Please refer to the Infiniband Switch User Guide for further details) • Changed DSCP configuration to be per port instead of global • Network Direct: Fixed race in NDK between handling of incoming connection and destruction of a listener • Network Direct: Fixed race between NDK object creation and usage • Improved TCB (Transmission Control Block) management on send • Improved transmit and receive in multi stream scenarios • Enabled hardware checksum offload for non TCP/UDP traffic with ConnectX3-Pro • Improved stability when handling OIDs during driver reset • Fixed performance tuning for 1GbE link • Fixed possible reset of driver during migration of large number of VMs at the same time • Fixed stability issues
RoCE	<ul style="list-style-type: none"> • Added RoCE IP based
ND	<ul style="list-style-type: none"> • Fixed wrong return value in <code>IND2Adapter::QueryAddressList</code>

Table 2 - New Features, Changes and Fixes in v4.80.50000

Category	Description
Infiniband	<ul style="list-style-type: none"> Added non-default PKey in VM
Performance	<ul style="list-style-type: none"> Optimized interrupt moderation values in SR-IOV VF mode for IPoIB Improved perf_tuning detection for the first port Improved performance in packet forwarding scenarios Decreased dropped packets rate for Ethernet significantly Changed default perf_tuning scenario to be “Balanced configuration” Various performance improvements
WMI/CIM	<ul style="list-style-type: none"> Added ability to read active RoCE configuration from hardware Added support for RoCE IP Based

2.2 New Features, Changes and Fixes in Version 4.70

2.2.1 New Features, Changes and Fixes in v4.70.50050 from v4.70.50040



This package version is 4.70.50050. The package contains the following versions of components.

- Bus, eth, IPoIB and mux drivers version is 4.70.10143
- The CIM provider version is 4.70.10143

Table 3 - New Features, Changes and Fixes in v4.70.50040

Category	Description
IPoIB	<ul style="list-style-type: none"> Fixed SM fail-over causing the driver to hang

2.2.2 New Features, Changes and Fixes in v4.70.50040 from v4.70.50000



This package version is 4.70.50040. The package contains the following versions of components.

- Bus, eth, IPoIB and mux drivers version is 4.70.10141
- The CIM provider version is 4.70.10141

Table 4 - New Features, Changes and Fixes in v4.70.50040

Category	Description
Generic	<ul style="list-style-type: none"> Optimized handling of “affinity change” on OID_RECEIVE_FILTER_QUEUE_PARAMETERS Added the ability to control the number of retries and timeout to check the device health before performing reset
Ethernet	<ul style="list-style-type: none"> Fixed missing pause response by sender when using DSCP/untag priority tag mode with ETS enabled

2.2.3 New Features, Changes and Fixes in v4.70.50000 from v4.61



This package version is 4.70.50000. The package contains mixed versions of components.

- Bus, eth, IPoIB and mux drivers version is 4.70.10126.
- The CIM provider version is 4.70.10130.

Table 5 - New Features, Changes and Fixes in v4.70.50000

Category	Description
Installation/Upgrade	<ul style="list-style-type: none"> • Fixed removal of virtual IPoIB ports in uninstallation • All user-space binaries are now signed • Fixed restoration process of DNS servers during upgrade • Fixed popping windows during installation/upgrade • Fixed missing 32 bit files in the catalog files
Generic	<ul style="list-style-type: none"> • Changed Ethernet and IPoIB event log messages to be more clear • Ported SDK project to Visual Studio 2013. • Fixed an issue which caused Mellanox miniport devices to be listed in “Devices and Printers” • Fixed Ethernet and IPoIB deadlock in power state change during shutdown/reboot • Fixed stability issues
IPoIB	<ul style="list-style-type: none"> • Added support for IPoIB SR-IOV Virtual Function (VF) over KVM Hypervisor (Beta level) • Added support for non-default pkey, as queried from OpenSM, on IPoIB SR-IOV VF over KVM. • Added IPoIB QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perfmon utility • Fixed part_man exit with return value 0 in case of error
Ethernet	<ul style="list-style-type: none"> • Added support for Ethernet SR-IOV over Windows Hyper-V Hypervisor (over Windows 2012 R2)^a • Added Virtual Ethernet Adapter support which enables using SMB Direct and HyperV (VMQ and NVGRE (over ConnectX®-3 Pro)) on the same port^b • Added lossless TCP buffer management when no receive WQE are available
RoCE	<ul style="list-style-type: none"> • Added ConnectX®-3 Pro support for RoCEv2 • Changed the transport name in vstat and ibstat to be RoCE v2.0 • Fixed ibstat behavior on devices with RoCE enabled • Fixed releasing of RDMA resources and reacquire them on power down and up. • Fixed RDMA Activity counters which didn't increase for ND traffic
ND	<ul style="list-style-type: none"> • Fixed hard-coded limitation of 4 SGEs
InfiniBand	<ul style="list-style-type: none"> • Fixed vstat printing of counters for Ethernet ports
Performance	<ul style="list-style-type: none"> • Improved perf_tuning setting in single CPU machines to avoid TX collision

a. Requires firmware v2.30.8000 and above

b. Requires firmware v2.31.5050 and above.

2.3 New Features, Changes and Fixes in v4.61 From v4.60



WinOF VPI version 4.61 was released as an intermediate release.

Table 6 - New Features, Changes and Fixes in v4.61

Category	Description
Installation/Upgrade	<ul style="list-style-type: none"> Fixed an issue preventing JumboPackets registry key to be restored correctly Ensured that uninstallation of Mellanox package in Virtual Machine leaves the system clean
Generic	<ul style="list-style-type: none"> Improved information in event log when a bad cable is detected Improved resiliency on error flow in Ethernet, IPoIB and bus drivers Fixed an issue which caused Mellanox devices to be listed in “Devices and Printers” and had “Safe Removal” UI
Performance	<ul style="list-style-type: none"> Added support OF IPv6 to all nd_*_* tests Enabled optimal interrupt moderation values in SR-IOV VF mode Stopped using NdisQueryNetBufferPhysicalCount to improve CPU utilization
IPoIB	<ul style="list-style-type: none"> Enabled searching for IBAT routes based on dest only instead of src,dest and added a mechanism preventing memory growth in IBAT entries Allowed any number of RSS processors, not only a power of 2 Ensured SR-IOV mode is not enabled for IPoIB ports, which resulted in confusing message in event log Fixed error statistics collection which could cause false error report Fixed a connectivity problem between Hyper-V VMs on the same host Fixed loopback issues in the virtualization environment Fixed stability issues
Ethernet	<ul style="list-style-type: none"> Added support for “unknown” link state indication Added support for DMA checks by driver verifier on SR-IOV Virtual Function Added support for NVGRE over LBFO Team Improved performance of handling change receive ring affinity request In SR-IOV mode, improved resiliency to driver failures in the Virtual Machine which could result in driver load failure in VM In SR-IOV mode, improved resilience in VF to PF communication Improved structure of INF file for SR-IOV Physical and Virtual Functions Fixed an issue that prevented receiving ARP traffic in NVGRE mode

2.4 New Features, Changes and Fixes in Version 4.60

2.4.1 New Features, Changes and Fixes in v4.60.17738 from v4.60.17718



This package version is 4.60.17738. The package contains the following versions of components:

- Bus and eth driver version 4.60.17718.
- The CIM provider version is 4.60.17718.
- The mux driver version is 4.60.17729.
- The IPoIB driver version is 4.60.17736.

Table 7 - New Features, Changes and Fixes in v4.60.17738

Category	Description
IPoIB	<ul style="list-style-type: none"> • Fixed using CQ after VMQ is closed • Fixed bad completion of VMQ QP that was caused by malformed WR
Ethernet	<ul style="list-style-type: none"> • LBFO: Fixed the team's MAC address uniqueness in the subnet of the team in Windows 2008 R2

2.4.2 New Features, Changes and Fixes in v4.60.17718 from v4.55



This package version is 4.60.17718. The package contains the following versions of components:

- Bus, eth, IPoIB and mux drivers version is 4.60.17718.
- The CIM provider version is 4.60.17718.

Table 8 - New Features, Changes and Fixes in v4.60.17718

Category	Description
Installation/Upgrade	<ul style="list-style-type: none"> • Enabled configuration changes saving upon Inbox and previous releases upgrade • Enabled CIM installation as a standalone package • Fixed an issue occurred when uninstalling and reinstalling the driver. The ConnectX-3 Pro Ethernet device was displayed in the Device Manager with a yellow bang (!). • Fixed an issues enabling the package's execution in modify mode resulting in driver being disabled
Generic	<ul style="list-style-type: none"> • Added support for a new report for bad cables^a • Fixed random parsing failures of string registry entries • Fixed compilation failure of "Hello_world" in the SDK • Fixed the return value of <code>ib_query_ca()</code> if failed to allocate resources for operation

Table 8 - New Features, Changes and Fixes in v4.60.17718

Category	Description
Performance	<ul style="list-style-type: none"> • Added support to IPv6-to-all nd_*_* tests • Fixed CPU utilization report in nd_*_* tests • Fixed correct bandwidth peak results in ibv_send_bw with UD QP • Fixed sync problems of bidirectional mode in ibv_read_bw/ibv_write_bw • Fixed an issue reporting incorrect adapter type in performance tuning log file
RoCE	<ul style="list-style-type: none"> • Fixed RoCE mode parsing
ND	<ul style="list-style-type: none"> • Added the ability to rearm a CQ in the kernel • Added the ability to handle LID changes • Changed connection timeout behavior. Added the STATUS_CONNECTION_REFUSED return value upon connection timeout. • Fixed missing completions when working with Completion Queue with single entry
IPoIB	<ul style="list-style-type: none"> • Added the ability to handle LID changes • Added support for iSCSI boot over IPoIB • Fixed unexpected behavior upon QP asynchronous event • Fixed bad completions of VMQ and NonVMQ modes in IPoIB • Fixed a failure occurred when setting the IPoIB adapter value to “SA Query Timeout” • Fixed propagation of the physical link disconnection to virtual (part_man) interface • Fixed BSOD caused by calling ib_join_mcast() with timeout_ms = 0 • Performance improvements in latency
Ethernet	<ul style="list-style-type: none"> • Added DSCP support over IPv4^a • Added traffic profile • Added IRQ dynamic moderation • Modified the CQ size to prevent CQ overrun • Changed the report link speed zero in case of disconnected network adapter • LBFO: Fixed port channel teaming with CISCO switch and Fabric Extenders traffic loose in Windows Server 2008 R2 • Fixed an issue related to packets sent with corrupted VLAN header when they were meant to be untagged • Fixed unexpected behavior upon QP asynchronous event • Fixed the ability to disable Wake on Lan (WoL) on NICs which supports it. • Stability fixes • Performance improvements

Table 8 - New Features, Changes and Fixes in v4.60.17718

Category	Description
WMI/CIM	<ul style="list-style-type: none"> • Added ControlledBy association to IBPort • Fixed ConformsToProfiles association for SoftwareIdentity and DriverIdentity • Fixed execution of all tests which were running when executing Diagnostic tests on one instance • Fixed a failure occurred when running MLNX_Card • Fixed the printing of diagnostics log • Fixed an issue preventing from get-event to show information after disabling the PCI device • Removed support for the following configuration: <ul style="list-style-type: none"> • ModeFlags • SingleMsixNum • MultiMsixNum • SingleEqNum • MultiEqNum • MaxContQuant • SlaveNum • DebugLevel • DebugFlags • UsePrio • NumFcExch • EnableQoS • BlockMcastLoopBack • InterruptFromFirstPacket • ProbeVf

a. Requires firmware v2.30.8000 and above.

2.5 New Features, Changes and Fixes in v4.55 From v4.40

Table 9 - New Features, Changes and Fixes in v4.55 (Sheet 1 of 2)

Category	Description
Generic	<ul style="list-style-type: none"> • Added support for Windows Server 2012 R2 Operating System • Added the ParentBusPath option to each port registry key • Added a new hardware ID for ConnectX®-3 Pro NICs • The QP numbers allocation is now round-robin manner • RecvCompletionMethod as Interrupt is no longer supported • Removed the LsoV1IPv4 from the registry/UI • Removed from the bus driver configuration the 'Non-DMA' option • Removed the TXRingNum option from the UI
NVGRE	<ul style="list-style-type: none"> • Added NVGRE hardware off-load support (for ConnectX-3 Pro cards only) • Added to the UI the *EncapsulatedPacketTaskOffload option when using ConnectX®-3 Pro NICs
Performance	<ul style="list-style-type: none"> • Added the nd_send_bw and nd_send_lat ND benchmarking tools • Fixed nd_*_bw to achieve better performance (memory buffer alignment) and consistent results

Table 9 - New Features, Changes and Fixes in v4.55 (Sheet 2 of 2)

Category	Description
Ethernet	<ul style="list-style-type: none"> Fixed the issue preventing messages to be sent in VLAN 0 when using many VMQ rings Added IP-IP checksum off-load support Added Ports TX arbitration/Bandwidth allocation per port The following ND providers, MLX4ND and MLX4ND2 are installed by default Fixed setting the correct SL in UD traffic over RoCE
InfiniBand	<ul style="list-style-type: none"> IPoIB performance improvements Fixed a part_man issue related to wrong statistics over virtual partman interfaces
RoCE	<ul style="list-style-type: none"> Enabled roce_mode value overwrite in case it exists during installation Fixed in ibv_devinfo the display of correct transport RoCE mode Added Sniffer for RoCE packets The used RoCE mode set upon driver load is printed into event log message

2.6 New Features, Changes and Fixes in v4.40 From v4.30

Table 10 - New Features, Changes and Fixes in v4.40 (Sheet 1 of 3)

Category	Description
Generic	<ul style="list-style-type: none"> Added a notification in the event log in case SMB is not supported in ConnectX@-2 firmware Added the trace tool for WPP tracing Added copyright to the SDK files Added WMI/Powershell support Fixed an issue causing the setup to fail upon perf_tuning failure during the installation. An error message will be printed in the installation log upon perf_tuning failure. Removed port setting registry key during uninstall Fixed an issue with the Mellanox adapter being shown on the USB removal menu, which caused the removal of the Mellanox adapter once removing the USB.
Performance	<ul style="list-style-type: none"> Set 512 RX buffers by default Removed TXRingNum Changed the perf_tuning setting to achieve a better performance tuning Added the nd_write_bw/nd_write_lat and nd_read_bw/nd_read_lat tools Fixed the perf_tuning indication of the last chosen tuning scenarios Fixed a crash in the ib_send_lat/bw utilities caused when the port link was down Fixed the “Restore to defaults” option in the perf_tuning tool. Now the default values are being restored

Table 10 - New Features, Changes and Fixes in v4.40 (Sheet 2 of 3)

Category	Description
Ethernet	<ul style="list-style-type: none"> • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perfmon utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive <p style="margin-left: 40px;">In addition to old values that are not supported anymore.</p> • Made mlx4_bus and Ethernet devices removable • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to mlx4nd2 • Fixed WoL support on NIC with a single port • Fixed the default RoCE configuration on NICs with a single ports • Fixed the values for the MTU and rate of the CM-REQ • Fixed miniport reset on sending scenarios • Removed the QoS attributes when disabling QoS • Enabled MaxRssProcessirs support of the following values: 1, 2, 4, 8, 16, 32, 64 • Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation • Network Direct: Fixed the swap of InboundReadLimit and OutboundReadLimit when creating an EndPoint and in Connector::GetConnectionData • Network Direct: Fixed disallowing creation of EndPoint with zero attributes in the Receive Queue • Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again • Network Direct: Fixed a failure when creating an EndPoint with zero attributes in the Receive Queue • Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perfmon utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive <p style="margin-left: 40px;">In addition to old values that are not supported anymore.</p>

Table 10 - New Features, Changes and Fixes in v4.40 (Sheet 3 of 3)

Category	Description
	<ul style="list-style-type: none"> • Made mlx4_bus and Ethernet devices removable • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to mlx4nd2 • Fixed WoL support on NIC with a single port • Fixed the default RoCE configuration on NICs with a single ports • Fixed the values for the MTU and rate of the CM-REQ • Fixed miniport reset on sending scenarios • Removed the QoS attributes when disabling QoS • Enabled MaxRssProcessors support of the following values: 1, 2, 4, 8, 16, 32, 64 • Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation • Network Direct: Fixed the swap of InboundReadLimit and OutboundReadLimit when creating an EndPoint and in Connector::GetConnectionData • Network Direct: Fixed disallowing creation of EndPoint with zero attributes in the Receive Queue • Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again • Network Direct: Fixed a failure when creating an EndPoint with zero attributes in the Receive Queue • Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host
InfiniBand	<ul style="list-style-type: none"> • On rare occasions, depends on the GUID assignment, the IPoIB MAC address can be assigned with a multicast MAC (the least significant bit of the most significant address octet is set to 1). In that case, all of the traffic over the IPoIB I/F is dropped. If you experience this issue, please contact Mellanox support. • Added active_mtu field to struct ib_port_attr_t • Added the option of vstat displaying the active_mtu of the ports • Allowed registration of a large Memory Region which is splitted to many segments • Fixed a bluescreen issue that occurred when disabling the interface after a TX stress over the VMQ • Fixed a failure of MPI/ND over InfiniBand • Added the option of ibv_devinfo displaying the correct MTU value after it was changed • Added the option of part_man printing the adapter name when the Port GUID is set to zero. • Added the option of part_man printing the leading zeroes of port GUID
Installation/Upgrade	<ul style="list-style-type: none"> • Prevented displaying a message to upgrade the firmware for OEM NICs if it has the latest firmware version • Removed portsetting registry key during uninstallation

2.7 New Features, Changes and Fixes in v4.30 From v4.20



WinOF VPI version 4.3 was released as an intermediate release.

Table 11 - New Features, Changes and Fixes in v4.3

Category	Description
Generic	<ul style="list-style-type: none"> Added support for a new provider called MLX4ND, which supports both NDv1 and NDv2 interfaces
Performance	<ul style="list-style-type: none"> Enabled performance tuning running according to the operating systems that are running over it. The keywords added to the registry in NDIS support Windows 2012 are: <ul style="list-style-type: none"> RssMaxProcNumber NumRSSQueues RSSProfile The keywords added to the registry in NDIS support Windows 2008 R2 are: <ul style="list-style-type: none"> RssBaseProcNumber MaxRssProcessors RssBaseProcGroup The rest of the keywords are added in all versions of NDIS. This change is based on: http://msdn.microsoft.com/en-us/library/windows/hardware/ff570864(v=vs.85).aspx
Ethernet	<ul style="list-style-type: none"> RoCE MTU value is no longer set to 1024 by default. All options stay as they are and can only be chosen if they were selected explicitly in the UI/registry. The current default state is as follows: The value is now derived from the MTU (or MaxFramSize, or Jumbo Packets value) and they are all aliases for the same value). The value is aligned to 256,512,1024,2048 in a way that it will be rounded down to the nearest power of two of the ETH MTU.
InfiniBand	<ul style="list-style-type: none"> Added ibdiagnet utility support

2.8 New Features, Changes and Fixes in v4.2 From v3.2.0

Table 12 - New Features, Changes and Fixes in v4.2 (Sheet 1 of 2)

Category	Description
Generic	<ul style="list-style-type: none"> • Modified RSS cores and changed VMQ affinity on the fly • Fixed restart issue when there are not enough MSI-X vectors for each machine core • Added support for K-GROUPS processors (more than 64 processors support) to allow assignment of MSI-X affinity for multiple processor groups. • Set an adequate number of MTTs to map all physical memory • Allocated firmware and ICM memory in chunks of non-paged memory instead of using contiguous physical memory. • Fixed RSS indirection table mapping building when there are less RX rings than RSS cores. • Fixed a bug, preventing standard work with BAR value more than 4GB. • Fixed memory leaks • Fixed error flows causing a Bluescreen in driver startup/unload • Fixed a Bluescreen occurrence upon shutdown due to leak in active resources • Changed device names in device manager and their hardware IDs. The changes were made to distinguish between ConnectX®-2 and ConnectX®-3: <ul style="list-style-type: none"> • for ConnectX-2: MLX4\ConnectX-2_Eth and IBA\ConnectX-2_IPoIB • for ConnectX-3: MLX4\ConnectX-3_Eth and IBA\ConnectX-3_IPoIB • Set QoS settings only for ConnectX-3. Changing the hardware ID, forces the OS to install new device and re-build the registry keys. • Added an event log to indicate driver failure upon start if there are two HCA burned with the same GUID. • Added firmware upgrade support as part of the setup process. The setup burns the new firmware only on Mellanox cards. Firmware burning failure does not prevent the driver's installation, therefore, it will show a warning. In this case, it is recommended to update the firmware manually. • Enabled configuration of TxRingNum registry key from the UI • Improved the "Port Protocol" dialog • Added Registry key documentation to the setup package
Performance	<ul style="list-style-type: none"> • Optimized code performance • Increased send parallelism • Memory used in receive flow is now allocated with the same affinity of the handling processor for faster access • Statistics parameters are now directly read from hardware instead of being calculated by software. • Added support for BlueFlame. BlueFlame is now the default working mode for all packets that have a descriptor which fits into a BF register (currently 256 bytes). Use "BlueFlame" registry key to enable/disable this feature. • Added support for RSS functionality on available processors numbers. Used to be restricted to start at the first processor. • Changed RSS registry defaults to give better out of the box performance • Added a performance UI to tune performance under various scenarios • Added a tool to tune performance under various scenarios

Table 12 - New Features, Changes and Fixes in v4.2 (Sheet 2 of 2)

Category	Description
Ethernet	<ul style="list-style-type: none"> • Added support for multiple TX rings • Added an option to verify that the number of multicast groups used is no higher than the firmware limits • Improved performance in virtualization when using VMQ
IPoIB	<ul style="list-style-type: none"> • Fixed a bug that prevented IBAL applications from working on machines with 2 cards • Fixed a bug that caused packet drop on remote node when the first sent packet is not an ARP or a multicast packet • Added support for multiple TX rings. The new driver advertises its TSS capability (using multiple TX rings). This feature will not be used when communicating with older version drivers (such as the inbox driver). This way IPv4 communication in a mixed environment that includes the inbox driver should work. IPv6 communication may be problematic in such an environment. • Eliminated some race conditions in IPoIB mcast handling • Added mcast addresses according to the Windows mcast table in addition to mcast addresses from igmp/mld queries, in order to connect to such mcasts that are not reported by the queries. • Removed redundant LID_CHANGE events. This fixes the redundant reconnection of all IPoIB mcasts. • Added support for VM migration. When a VM is migrated from one host to another it does not send a new ARP to hosts it already knew when it was sitting on the first host. In order to prevent communication interruption to such hosts, the driver will initiate an ARP to learn them again. • Fixed IPoIB VMQ affinity update bugs • Fixed IPoIB VMQ parent queue management race during reset • Fixed a bug in passing DHCP packets from a Linux VM • Fixed bug in RDMA statistics for NDK • Removed affinity restriction to group 0 for NDK • Added support for the part_man tool. It can create only one additional IPoIB interface per port GUID. This feature can be utilized to use different IPoIB interfaces for SMB and Hyper-V. The virtual ports configuration is removed during uninstall • Installed ND by default • The following tools were added to the setup package: <ul style="list-style-type: none"> • ibv_asyncwatch.exe • ibv_atomic_bw.exe • ibv_atomic_lat.exe • ibv_devinfo.exe • ibv_rc_pingpong.exe • ibv_rdma_bw.exe • ibv_rdma_lat.exe • ibaddr.exe • ibcacheedit.exe • iblinkinfo.exe • ibqueryerrors.exe • ibsysstat.exe • saquery.exe • smpdump.exe

3 Features/Support to be Deprecated in the Next Release

The following support will be deprecated in the next release:

- ConnectX®-2 adapter cards

4 Beta Features

Table 13 - Beta Features

Category	Description
WinVerbs	WinVerbs is currently at beta level.
ibdump	ibdump is currently at beta level.
IPoIB	IPv6 support of IPoIB in an SR-IOV guest OS over KVM is at beta level.
	Multiple PKeys is at beta level

5 Unsupported Functionality/Features

The following are the unsupported functionality/features in WinOF Rev 4.80.50000:

- ND over WinVerbs provider
- SRP

5.1 ConnectX®-2 Adapter Limitations

- Limited NDK support - for details contact Mellanox Support.
- CQ to EQ feature is not supported - consequently, Multi-core performance may be non optimal
- Mellanox perfmon counters are not supported
- Auto Sensing of link protocol (InfiniBand or Ethernet) is not supported
- ibdump is not supported
- QoS is not supported

6 Known Issues

6.1 Generic Issues

Table 14 - Generic Issues

Issue	Workaround
Pinning all the physical memory (used by RDMA operations, such as register memory, pin user memory) on the machine, on Operating Systems prior to Windows Server 2012, may cause the machine to hang.	Avoid pinning the whole machine memory in those Operating Systems.
When the tracer tool traces many events, it may consume a large amount of memory (up to several GB RAM).	Reduce the verbosity level.
The displayed MAC address in the DHCP server of virtual IPoIB I/F may display a wrong data (FF00.....) although the I/F is still fully functional.	-
When running applications that use ND or libibumad (such as OpenSM) the system might get to an unstable state when trying to shutdown/restart/hibernate it.	Close all applications that use ND or libibumad before performing shutdown/restart/hibernate.
Activating NC-SI in WinOF v0.52 may cause driver's loading failure when using an older firmware version than 2.30.8000.	Do not enable NC-SI in machines that WinOF v0.52 is installed in.
The maximum values returned by the <code>ib_query_ca()</code> function (for example: <code>max_qp</code> , <code>max_mr</code>) are the upper limits of the supported resources by the device. However, it may be impossible to use these maximum values, since the actual number of any resource that can be created may be limited by the machine configuration, the amount of host memory, user permissions, and the amount of resources already in use by other users/processes.	-
Running Ntttcp without the "-a X" flag (X > 1) in a NIC configured with 10GbE, may cause low bandwidth in TCP single stream.	Run Ntttcp with "-a 8" for best performance
Active links disappear after changing the cable connectivity from Ethernet to InfiniBand or vice versa.	Disable and enable the <code>mlx4_bus</code> interface from the Device Manager.

Table 14 - Generic Issues

Issue	Workaround
<p>On ConnectX®-2/ConnectX®-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/driver utilities that read the GUID via device firmware (e.g., using <code>ibstat</code>). <code>Mlxburn/flint</code> return <code>0xffff</code> as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.</p>	<p>Please use the GUID value returned by the fabric/driver utilities (not <code>0xffff</code>).</p>
<p>Changing the default ND providers may cause random errors, such as: <code>0xC0000238</code> <code>(NT_STATUS_ADDRESS_ALREADY_ASSOCIATED)</code> on <code>Connect()</code> or with <code>0xC0000236</code> <code>(NT_STATUS_CONNECTION_REFUSED)</code> on <code>Accept()</code>. These errors can be safely ignored.</p>	<p>-</p>
<p>The UI displays the network interface's state incorrectly even after disabling and re-enabling the "disconnected" network interfaces.</p>	<p>To see the correct state and link speed, perform one of the following:</p> <ul style="list-style-type: none"> • Run <code>Get-netadapter Powershell cmdlet</code> or • Right click on that interface from "Network Connections" and click on status
<p>SR-IOV should be enabled only when both ports are Ethernet. If SR-IOV is enabled when both ports are in auto sensing mode, the ports type is forced to be Ethernet.</p>	<p>-</p>
<p>WoL is not supported on Windows 2008 R2.</p>	<p>-</p>
<p>Performing the following scenario causes memory leak of 64 bytes on the Virtual Machine:</p> <ul style="list-style-type: none"> • Creating VFs • Restarting VM associated to VF(s) • Disabling MLNX driver on the Hypervisor 	<p>Do not disable the MLNX driver on the Hypervisor</p>
<p>iSCSI boot over Windows 2008 R2 is not supported</p>	<p>-</p>
<p>Clearing the Mellanox device counters through <code>perfmon</code> does not always work.</p>	<p>Restart the driver</p>
<p>When running in SRIOV mode and Hyper-V, the same driver version should be used on both the Hypervisor and the virtual machine.</p>	<p>-</p>

6.2 InfiniBand Issues

Table 15 - InfiniBand Issues

Issue	Workaround
InfiniBand application that was compiled with an SDK version earlier than WinOF v4.40 is not binary compatible.	Recompile InfiniBand application with WinOF v4.40 and above. ND application is backward compatible and older applications over ND do not have to be recompiled.
WMI does not work due to lack of permissions.	Change the execution policy. Run: <code>Set-ExecutionPolicy AllSigned</code>
<p>WinOF 4.40 and above IPoIB driver is not IPv6 compatible with earlier driver versions, including Windows 8 Inbox driver.</p> <p>If WinOF 4.50 IPoIB node receives an icmpv6 message from the legacy IPoIB node, the following event will appear in the event log:</p> <p>"IPoIB driver detected a possible incompatibility with Windows 8 inbox IPv6 support due to which there is no IPv6 connectivity".</p>	<p>To enable compatibility mode, add:</p> <ul style="list-style-type: none"> Win8InboxCompatibilityMode REG_SZ registry key with the value of 1 in the IPoIB interface registry. <p>Note: All IPoIB nodes must use the same mode to enable IPv6 compatibility with earlier driver versions.</p> <p>We recommend upgrading all hosts to the new driver version from http://mellanox.com or use Windows 8 compatibility mode</p>
Functionality such as CQ to EQ; Auto Sensing, QoS are not supported in ConnectX@-2 adapter cards.	-
On rare occasions, as a result of port configuration change (IB/ETH) the UI may get stuck for up to a few minutes. This effect does not require any user action. The UI returns to its proper functionality after a few minutes.	-
WinOF Inbox driver does not support upgrade. When installing WinOF v4.40 and above on a Windows Server 2012 and above machine, the Inbox driver is uninstalled prior to starting the new installation and any previous configurations is lost. The Inbox driver will be reinstalled automatically when the new driver is uninstalled.	-
Without separate ports for each stream, WinSock multiplexes every packet to every subscriber socket and then filters it out.	Use different UDP ports to get higher performance when using multicast packets.
A virtual IPoIB interface, created by the <code>part_man</code> utility, reports an Active state when the physical link is in the Initializing state and OpenSM is not running in the subnet	-
The "Packets Received Discarded" and "Packets Received Errors" counter may display wrong results.	-

Table 15 - InfiniBand Issues

Issue	Workaround
Firmware upgrade may fail during installation if there was a prior firmware upgrade on the machine without a reboot after it. A firmware upgrade failure does not fail the whole installation.	Upgrade the firmware manually.
The drivers' functionality is limited up to 128 cores.	-
Connection failure on ND tests while machine A have IBAL provider and machine B have MLX4ND provider.	-
Hibernate and Sleep are not functional when user-space is using its resources.	-
Calling <code>ib_join_mcast()</code> with <code>timeout_ms = 0</code> may result in a BSOD.	Avoid calling this function with <code>timeout_ms = 0</code> .
IPoIB does not support: <ul style="list-style-type: none"> • MAC address change • QoS (packet priority) • Load balancing and fail-over (LBFO) • Connected Mode • Partition 	-
Memory registration on 32bit machines is limited to up to 256GB.	-
In an interoperability environment that has both Linux and Windows OSs, the MTU value must be the same, otherwise packets larger than the minimum will not go through. The default MTU for Linux is 2K and for Windows is 4K.	-
OpenSM does not run as a service during installation since the SM requires the GUID parameter to decide which port to work on. Setting it on setup causes it to work only on the first port and not the others.	To run OpenSM as a service, assuming the package was installed in the default path, use: <pre>sc create opensm binPath= "c:\Program Files\Mellanox\MLNX_VPI\IB\Tools\opensm.exe"</pre> To start the service run: <pre>sc start opensm</pre>
Tools issues: <ul style="list-style-type: none"> • <code>ibportstate</code> does not work on RoCE ports • <code>ibdiagpath</code> may crash on Hyper-V machines 	-
If an application which uses InfiniBand runs while the driver is being restarted, a bluescreen or an NMI may occur.	Stop all InfiniBand applications including OpenSM upon driver restart.
If OpenSM is up during driver restart on the same machine, it might stop working when the driver is being backed up.	To resume operation, stop OpenSM and restart the driver while OpenSM is down.

Table 15 - InfiniBand Issues

Issue	Workaround
Sometimes during traffic, the latency of the IPoIB can get higher (this behavior is inconsistent).	Set the following registry in the IPoIB network interface: RecvIntModTime=0 Please note, this action increases the CPU utilization.
No communication between the physical interface and a VM that uses vSwitch created over virtual IPoIB (Pkey), and vice versa.	-
Calling <code>ib_query_cq()</code> on a CQ which was created with 1 entry indicates that there are 0 entries in the CQ.	-
An SR-IOV Virtual Machine is loaded in InfiniBand mode when no OpenSM is available in the subnet. The following event appears in the event log: “detected a null port GUID for port <NUMBER>. A Virtual Function device may have a null port GUID if there is no OpenSM instance on its network. Please make sure the network has an active OpenSM and restart the driver.”	Start the OpenSM and restart the driver.
For InfiniBand SRIOV guest, OpenSM Assigned GUIDs are not supported and may cause unexpected behavior	Work only with Administrator assigned GUIDs.
For InfiniBand SRIOV guest, ND traffic does not work on an SRIOV adapter when a Para-Virtualization adapter configured on the same virtual machine and the same subnet as IPoIB.	Disable the Para-Virtualization adapter.

6.3 Ethernet Issues

Table 16 - Ethernet Issues

Issue	Workaround
Disabling the “Priority & VLAN tag” in the UI which VLANID is configured, may result in sending packets with the configured VLANID.	Remove the VLANID before disabling the “Priority & VLAN tag”.
When working with LBFO, the teamed interface disappears after machine reboot. The issue applies to Windows Server 2012 R2	Delete the existing teamed interface and create a new one.
WakeOnMagicPacket registry key is not added to the registry although WoL is supported by the driver and by the NIC.	-
When the ports of the device are configured as Ethernet only, <code>ibstat/vstat</code> may display wrong information.	-

Table 16 - Ethernet Issues

Issue	Workaround
High multicast drop rate on multicast storming.	Use “Multicast traffic” tuning option under the performance tab. For further information, please refer to section “Tunable Performance Parameters” in the User Manual.
Driver installation requires deletion of the <code>m1x4_bus.sys</code> file in Windows Server 2008 R2 and WLH OSs when using the PXE package.	Delete the <code>m1x4_bus.sys</code> file and reboot the machine to install the driver.
When there is a stress in TCP connection establishments, some of those connections may fail.	Increase the Ring queue sizes: <ul style="list-style-type: none"> • ReceiveBuffers - controls the receive ring size • TransmitBuffers - controls the transmit ring size
The DCB component specifies a default traffic classification that is applied to all egress packets that do not match other classification conditions. In this case, the network adapter assigns the IEEE 802.1p priority level that is associated with the default classification to these egress packets. The default traffic classification has the following attributes: <ul style="list-style-type: none"> • It has a traffic classification condition of type <code>NDIS_QOS_CONDITION_DEFAULT</code>. • It is the first traffic classification defined in the array of <code>NDIS_QOS_CLASSIFICATION_ELEMENT</code> structures. 	-
RDMA Activity counters do not count during NetworkDirect RoCE traffic.	-
In SR-IOV mode, all ports are in pure Ethernet mode and RDMA is not supported on either port.	-
In SR-IOV mode, enabling SR-IOV in the BIOS may change the interfaces names. If any VSwitch is bounded to an interface whose name was changed, there will not be any way to perform any operation on them.	Unbond all Vswitches from the NIC's interfaces before enabling SR-IOV in the BIOS.
GRE traffic steering by inner MAC and by outer MAC simultaneously is currently not supported.	Configure steering or by inner MAC, or by outer MAC.
If VMQ set filter requests are accepted without a GRE flag (i.e. requested steering by outer MAC), the GRE packets do not reach that VMQ.	Set the bus driver registry <code>keyAcceptGREbyOuterMAC_P1/2</code> per port to accept GRE traffic by outer MAC and to duplicate L2 steering rule to <code>L2_TUNNELING</code> rule for each VMQ set filter request without GRE flag. Note: For regular NVGRE Hyper-V scenarios the value of the registry key below must be set to <code>0:AcceptGREbyOuterMAC_P1/2</code>

Table 16 - Ethernet Issues

Issue	Workaround
Disabling the ethernet adapter with more than 25 VLANs configured over windows 2008 R2 \ windows 7, may result with a non-responding server.	Reduce the number of configured VLANs prior to disabling the ethernet adapter.
Virtual Ethernet Interfaces created by <code>vea_man</code> are not tuned by the automatic performance tuning script.	For optimal performance need to follow the performance tuning guide and apply relevant changes to the VEA interface
In SR-IOV, 40Gbps interfaces of vSwitch interface on HyperV, and the VMNIC are identified as 10Gbps. Despite the incorrect speed identification, they still achieve 40Gbps performance.	-

6.4 Quality of Service Issues

Table 17 - Quality of Service Issues

Issue	Workaround
Running Quality of Service (QoS) commands without the parameter <code>"-PolicyStore ActiveStore"</code> may cause machines to load without Quality of Service policy.	Store the QoS policy in the ActiveStore

6.5 RoCE NVGRE Issues

Table 18 - RoCE NVGRE Issues

Issue	Workaround
RoCE does not support: <ul style="list-style-type: none"> Traffic cannot go through the router. It works in the same subnet only Multicast traffic VLAN Layer 3 feature 	-
In machines with heterogeneous NICs: a NIC which supports RoCE v2, and a NIC which does not support RoCE v2 the following issues might arise: <ul style="list-style-type: none"> Connect-X®-3 Pro is loaded with the transport type RoCE v2 Connect-X®-3 is loaded with the transport type RoCE 	-
When NVGRE off-load is enabled, the GRE traffic cannot be accepted as a regular L2 traffic and requires special <code>L2_TUNNELING</code> steering rules. In such case the GRE packets are dropped or directed to promiscuous queue.	-

Table 18 - RoCE NVGRE Issues

Issue	Workaround
Using different versions of RoCE in your cluster is not supported.	Use the same RoCE version in all the cluster in the Ethernet ports.

6.6 Performance Issues

Table 19 - Performance Issues

Issue	Workaround
When using WinOF 4.40 or above, low throughput will be seen on 40GbE adapters when QoS is enabled.	Disable QoS when it is not in use: <ul style="list-style-type: none"> Open a PowerShell prompt. Run: <pre>Disable-NetAdapterQos -name <Interface Name></pre> where <Interface Name> is e.g. "Ethernet 1"
<code>perf_tuning</code> is supported only when one of the two NUMA nodes are in use.	-
Running performance benchmarks for a short period of time (< 1 sec) may provide bad latency in IPoIB and Ethernet.	Set "Rx Interrupt Moderation Profile" and "Tx Interrupt Moderation Profile", to "Low Latency" to avoid bad latency. Note: This may increase CPU utilization.
The driver uses optimal interrupt moderation values for 10 GbE SR-IOV VF scenario. For other scenarios, the optimal values yet to be found.	-
In WinOF 4.80, when using the performance tuning tool on a single port device, the tool crashes.	Manually configure the desired parameters according to the Performance Tuning Guide.

6.7 IPoIB Issues

Table 20 - IPoIB Issues

Issue	Workaround
In VMQ and IPoIB, when performing VM Migration or modifying VMQ configuration during traffic with a large number of active VMs, the system may experience temporary network connectivity issues or failure of the VM migration	-
When using non-default Pkey for IPoIB in an SR-IOV guest OS, using the same partition with different membership might cause connectivity problems.	Work only with full membership PKeys as the non-default Pkey.
vIPoIB ports which are created via the <code>part_man</code> tool over the same partition with different membership, might cause connectivity problems.	Use <code>part_man</code> with the full membership value (the partitioning will be enforced by SM configuration).
IPoIB interfaces' teaming is currently not supported.	-

6.8 Hyper-V Issues

Table 21 - Hyper-V Issues

Issue	Workaround
<p>When the vSwitch is detached from the ETH\IPoIB device while the driver is disabled, the device does not reacquire the static IP it had before the attachment of the vSwitch. When the vSwitch is attached to the ETH\IPoIB device while there is no link, it will not receive the device IP when the link is back up.</p>	-
<p>After attaching the vSwitch to the ETH\IPoIB device, changing the “Jumbo Packet” registry key on the ETH\IPoIB device does not affect the vSwitch configuration and vice versa. For example, if the user sets the “Jumbo Packet” on the ETH\IPoIB device to X, and the “Jumbo Packet” on the vSwitch to X+Y, X+Y sized packets will be passed from NDIS down to the driver and they will be dropped by it.</p>	Reattach the vSwitch to sync with the value set in the ETH\IPoIB device.
<p>Unexpected behavior might occur when running in a virtualized environment and creating two virtual switches bound to each of the ports of a dual port NIC and then using both of them with two vNICs from the same VM.</p>	-
<p>When moving an IPoIB interface in a VM from non-VMQ to VMQ or from VMQ to non-VMQ, a reset to the NIC may occur and in the event log the following message will appear: <code>`device reports a "CQE error" on cqn <number> qpn <number> Status <number>.</code> Therefore, the HCA Nic will be reset. (The issue is reported in Function <function>). For more information refer to details.'</p>	-
<p>In IPoIB when using long Multicast traffic from a Virtual Machine (VM) to an external host there might be up to 0.5% loss in 5% bursts</p>	-
<p>Hyper-V is at low bandwidth on LBFO vSwitch, Windows 2012</p>	-
<p>In Ethernet to achieve better iperf TCP performance between a Linux VM and a Windows VM on different hosts, when using MS MUX over the Ethernet driver, use the non VMQ mode for the VMs.</p>	-

Table 21 - Hyper-V Issues

Issue	Workaround
After disabling and enabling a port on a guest, a ping to it may be renewed after a minute. The ARP requests sent by Windows are less frequent as the time passes. If the guest port was down for a while, it could take time until Windows decides to send another ARP request to it.	-
When VMQ is enabled after reset, the driver loads all the VMQs that existed before the reset. However, it is not guaranteed that each VMQ will receive the same QP number it had before the reset. This can cause some delay as a result of resetting before connectivity is reestablished. The delay is caused by the time it takes for the ARP table to update after initiating the Gratuitous ARP.	-
The IPoIB non-VMQ mode is supported only when the VMQ is enabled according to the registry values.	To use the non-VMQ mode for a VM, change its settings as follow: <ul style="list-style-type: none"> • Press “Settings” on the VM • Go to Network Adapter -> Hardware Acceleration • Un-check the “Enable virtual machine queue”

6.9 Installation / Upgrade Issues

Table 22 - Installation / Upgrade Issues

Issue	Workaround
Upgrading the driver while the UI is opened with the “ConnectX NIC device” may cause the installation process to never end.	Close the UI before driver upgrade.
Rebooting the machine reboot while uninstalling WinOF, may result in installation failure.	Delete Mellanox components from HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\DIFx-App\Components. The Mellanox components are mlx4eth63, ipoib6x and mlx4_bus
Canceling the installation process may leave the bus driver in a disable state. The driver appears in a yellow bang containing the following error message: “Windows cannot start this hardware device because its configuration information (in the registry) is incomplete or damaged. (Code 19)”.	Scan for new hardware and reboot the machine.
Downgrade is not supported.	Uninstall the current version and install the older one.

Table 22 - Installation / Upgrade Issues

Issue	Workaround
<p>Occasionally, upon driver upgrade due to mismatch of driver versions, the following error message is displayed:</p> <p>"There is an interface mismatch between ETH driver and the bus driver. The ETH driver interface version is 50 while the bus driver interface version is 17. As a result the ETH driver has failed to start.</p> <p>This happened due to setup failures or partial update of the drivers. In order to resolve the issue, please reboot the computer"</p> <p>The message can be safely ignored.</p>	-
<p>Configuration is not restored when replacing a ConnectX-2/ConnectX-3 NIC with a ConnectX-3 Pro NIC located on the same PCI slot</p>	Clean the old network adapter configuration prior to upgrade.
<p>Configuration can be restored only in Windows Server 2012 and above</p>	-
<p>IPv6 configuration restore is not supported</p>	-
<p>Upon upgrade, the following Registry Key values will be overwritten with the following:</p> <ul style="list-style-type: none"> • *ReceiveBuffers = 512 • *MaxRssProcessors = 8 • *RssBaseProcNumber = 0 • *NumRSSQueues = 8 • *RssMaxProcNumber = 63 • *RssProfile = 1 • DefaultRecvRingProcessor = -1 • TxInterruptProcessor = -1 • TxForwardingProcessor = -1 • RxIntModerationProfile = 1 • TxIntModerationProfile = 1 • RecvCompletionMethod = 1 • SingleStream = 0 • TxRingNum = 8 	-
<p>Upon upgrade the following Ethernet Registry Keys will be deleted:</p> <ul style="list-style-type: none"> • SendCompletionMethod • UseRSSForRawIP • UseRSSForUDP 	-
<p>Upon upgrade the SendCompletionMethod IPoIB Registry Key value will be modified as follow:</p> <ul style="list-style-type: none"> • SendCompletionMethod = 0 	-

Table 22 - Installation / Upgrade Issues

Issue	Workaround
Upon upgrade the following IPoIB Registry Keys will be deleted: <ul style="list-style-type: none"> • UseRSSForRawIP • UseRSSForUDP 	-
Uninstalling the driver on Windows Server 2008 R2 with LBFO configuration results in the appearance of a pop-up window requesting to close several running applications.	Choose “Do not close applications”. This action allows the uninstallation of the driver. A Reboot may be required. Rebooting the server before uninstalling the driver when LBFO is configured will eliminate this pop-up completely.
Running a downgrade in silent mode is not supported. Upon downgrade the return code will always be 0.	-
Uninstalling the driver after upgrade won't remove the directory %ProgramFiles%\Mellanox	-
Uninstalling the driver when multiple of VLANs are configured never ends.	Remove the VLANs before uninstallation.
The installation process does not close any applications running in the background, and may cause a BSOD as a result of a stuck cmd.	It is recommended to close all running applications prior to upgrading the driver.
Installation/upgrade fails due to PNP failure to copy the driver files to the driver store, and the following text is printed in the event logs: Fault bucket, type 0 Event Name: PnPDriverImportError Response: Not available Attached files: C:\Users\ <user>\App-Data\Local\Temp\DMI151A.tmp.log.xml C:\Program Files\Mellanox\MLNX_VPI\ETH\mlx4eth63.inf </user>	Reboot the machine and reinstall
Installation/upgrade fails due to failure to stop the WMI service, and the following text is printed in the installation log: “CustomAction StopWMIService returned actual error code 1603”	Kill the WMIPrvSE.exe tasks in the task manager and reinstall.

Table 22 - Installation / Upgrade Issues

Issue	Workaround
<p>Following the upgrade of Mellanox driver to WinOF-4.60 and above or on servers with no internet access, the first PowerShell command might be stuck for ~2-3 minutes before its completed.</p> <p>According to the following content, this issue is related to .Net framework version or an issue with the internet access: http://www.minasi.com/forum/topic.asp?TOPIC_ID=39253</p>	<p>Run the following script on the server to optimizes loading PowerShell DLLs:</p> <pre>\$Env:PATH = [Runtime.InteropServices]::GetRuntimeDirectory() [AppDomain]::CurrentDomain.GetAssemblies() % { \$pt = \$_.Location if (! \$pt) {continue} if (\$cn++) {''} \$na = Split-Path -Leaf \$pt Write-Host -ForegroundColor Yellow "NGENing \$na" ngen install \$pt }</pre>
<p>On ConnectX-3 cards only, when upgrading from Windows 2012R2 Inbox driver to WinOF, the RoCE mode setting in the registry is not properly transferred to the new driver. In case a non-default value was used it will not be configured following the upgrade.</p>	<p>Reconfigure the RoCE Mode setting manually.</p>
<p>Uninstall from the Device Manager is currently not supported</p>	<p>-</p>

6.10 Utilities Issues

Table 23 - Utilities Issues

Issue	Workaround
<p>ibdump may encounter packet drops upon a burst of more than 4096 (or 2^max-burst) packets.</p>	<p>-</p>
<p>Packets loss is not reported by ibdump.</p>	<p>-</p>
<p>Running ibdump on a RoCE Ethernet port may decrease the functional bandwidth due to the overhead of creating extra copy for each packet. This may lead to packet drops on the link.</p>	<p>Verify Ethernet flow control is enable to ensure a lossless link</p>
<p>Pcap file issues for RoCE IP Based:</p> <ul style="list-style-type: none"> The packets 'capture-time' field is wrong (may be 0 or negative numbers). For every captured packet, an additional 0 size flame is added. This appears in Wireshark as a 'malformed Ethernet packet'. 	<p>-</p>
<p>Sniffing over IB ports is currently not supported</p>	<p>-</p>

6.11 CIM/WMI Issues

Table 24 - CIM/WMI Issues

Issue	Workaround
Running Microsoft CIM cmdlets operations and their derived classes on classes <code>MLNX_NetAdapterSettingData</code> and <code>MLNX_NetAdapterRoceSettingData</code> is not supported. Calling those commands may cause the debugger, if connected to the machine, to assert.	Use <code>DriverCoreSettings</code> instead.
For PCI Gen3, <code>PcieLinkSpeed</code> is reported as "Unknown" when running <code>Get-NetAdapterHardwareInfo</code> Powershell cmdlet	-

7 API Changes

7.1 API Changes in WinOF Rev 4.80.50000

The following are the API changes in WinOF Rev 4.80.50000

Table 26 - API Changes in WinOF Rev 4.80

Name	Description
RDMA_TRANSPORT_RDMAOE_1	It is an alias to: RDMA_TRANSPORT_RDMAOE
RDMA_TRANSPORT_RDMAOE_1_25	Added enumerated values
is_roce(), is_mac_based_roce(), is_ip_based_roce(), is_rroce_or_ip_based_roce()	Added new functions
struct ib_wc_t	p_next was replaced with an anonymous union which contains two fields: p_next and qp_context

7.2 API Changes in WinOF Rev 4.70

The following are the API changes in WinOF Rev 4.70

Table 27 - API Changes in WinOF Rev 4.70

Name	Description
ib_get_port_spl_qp()	Added a new function
ib_get_mad_inner()	Changed API (one more input parameter was added)
ib_get_mad()	Changed API (one more input parameter was added)
VERBS_MINOR_VER	Increased its value, 000a -> 000c
UNBOUND_PORT_NUM	Added a new macro

7.3 API Changes in WinOF Rev 4.60

The following are the API changes in WinOF Rev 4.60:

Table 28 - API Changes in WinOF Rev 4.60

Name	Description
IB_MOD_QP_CHANGE_COUNTER_INDEX	Added a new macro
struct ib_qp_mod_t	Added the field state.rtr.counter_index

7.4 API Changes in WinOF Rev 4.55

The following are the API changes in WinOF Rev 4.55:

Table 29 - API Changes in WinOF Rev 4.55

Name	Description
RDMA_TRANSPORT_RRDMAOE_1_5	Added enumerated values
RDMA_TRANSPORT_RRDMAOE_2_0	Added enumerated values
RDMA_TRANSPORT_RRDMAOE	It is an alias to: RDMA_TRANSPORT_RRDMAOE_1_5
is_rroce(), is_xroce()	Added new functions
IB_AC_SNIFFER	Added a new macro
struct ib_qp_mod_t	Added the field state.init.flags

7.5 API Changes in WinOF Rev 4.40

The following are the API changes in WinOF Rev 4.40:

Table 30 - API Changes in WinOF Rev 4.40

Name	Description
VERBS_MINOR_VER	Increased its value, 0009 -> 000a
enum eth_link_speeds	Added enumerated values
struct ib_port_attr_t	<ul style="list-style-type: none"> The mtu field was separated into two fields: <ul style="list-style-type: none"> max_mtu (maximum MTU supported by the port) active_mtu (actual MTU which the port is configured with) Added the eth_link_speed field
WR_SEND_INV	Added enumerated values
struct ib_send_wr_t	The type of invalidate_rkey was changed from net32_t -> ib_net32_t
IB_SEND_OPT_SKIP_DOORBELL	Added the send Write flag