



Mellanox WinOF VPI

Release Notes

Rev 5.35

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

© Copyright 2016. Mellanox Technologies Ltd . All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies , Ltd.

All other trademarks are property of their respective owners .

For the most updated list of Mellanox trademarks , visit <http://www.mellanox.com/page/trademarks>

Table of Contents

Release Update History	5
Chapter 1 Overview.....	6
1.1 WinOF VPI Package Contents	6
1.2 WinOF VPI Package Contents for Nano Server	7
1.3 Supported Operating System Versions	7
1.4 Supported Network Adapter Cards	8
1.4.1 Firmware Versions	9
1.5 RoCE Modes Matrix	9
Chapter 2 Changes and New Features in Rev 5.35.....	10
2.1 Beta Features	10
2.2 Unsupported Functionality/Features	11
Chapter 3 Known Issues	12
Chapter 4 Bug Fixes History.....	37
Chapter 5 Change Log History.....	54

List of Tables

Table 1:	Supported Operating Systems	7
Table 2:	Supported Network Adapter Cards	8
Table 3:	Firmware Versions	9
Table 4:	RoCE Modes Matrix	9
Table 5:	RoCE v2 UDP Port Matrix	9
Table 6:	Changes and New Features in Rev 5.35	10
Table 7:	Beta Features	10
Table 8:	Known Issues.	12
Table 9:	Fixed Bugs List	37
Table 10:	Change Log History	54

Release Update History

Release	Date	Description
Rev 5.35 - FUR 2	March 21, 2017	<ul style="list-style-type: none"> Package version 5.35.52000 Driver version 5.35.12978 Added items 991194 and 995925 to “Bug Fixes History” on page 37.
Rev 5.35 - FUR 1	February 9, 2017	<ul style="list-style-type: none"> Package version: 5.35.51100 Driver version: 5.35.12970 Removed entries #966761 and #964639 from “Known Issues” on page 12 and added them to “Bug Fixes History” on page 37.
Rev 5.35	February 1, 2017	<ul style="list-style-type: none"> Package version: 5.35.50000. Driver version: 5.35.12965 Initial release of this WinOF version.

1 Overview

These are the release notes for the second fast update release of Mellanox WinOF VPI Drivers Rev 5.35 for Windows, supporting Mellanox ConnectX®-3 and ConnectX®-3 Pro network adapters.

Mellanox WinOF is composed of several software modules that contain InfiniBand and Ethernet drivers and utilities for ConnectX®-3 and ConnectX®-3 Pro adapter cards.



Windows Server 2012, Windows Server 2012 R2 and Windows Server 2016 include WinOF Inbox drivers which are a subset of the full WinOF VPI software package. As the Inbox drivers do not provide the full performance and functionality available with the WinOF VPI package, we recommend installing the full and latest WinOF VPI package.

1.1 WinOF VPI Package Contents

The Mellanox WinOF Rev 5.35 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)
 - Network Direct (ND)
 - Network Direct Kernel (NDK) Provider Interface
- Mellanox Ethernet LBFO driver for Windows Server 2008 R2 Mellanox IPoIB failover driver
- 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 support¹
- Software Development Kit (SDK)
- Documentation

1. Supported in Windows Server 2012 and above, and Windows Client 8.1 and above.

1.2 WinOF VPI Package Contents for Nano Server

In the package of Mellanox WinOF Rev 5.35.51000 for Windows, the following components are supported over Nano server:

- Core and ULPs:
- IB HCA low-level drivers (mlx4)
- IB Access Layer (IBAL)
- Ethernet driver (ETH)
- IP over InfiniBand (IPoIB)

1.3 Supported Operating System Versions

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

Table 1 - Supported Operating Systems

Virtualization Mode	Supported Host OS	Supported Guest OS
Native (non-virtualization)	Windows Server 2008 R2 ^a	N/A
	Windows Server 2012	N/A
	Windows Server 2012 R2	N/A
	Windows Server 2016	N/A
	Windows 7 Client (64 bit only)	N/A
	Windows 8.1 Client (64 bit only)	N/A
	Windows 10 Client (64 bit only)	N/A
Hyper-V (non-SR-IOV)	Windows Server 2008 R2	<ul style="list-style-type: none"> • Windows Server 2008 R2
	Windows Server 2012	<ul style="list-style-type: none"> • Windows Server 2008 R2 • Windows Server 2012
	Windows Server 2012 R2	<ul style="list-style-type: none"> • Windows Server 2008 R2 • Windows Server 2012 • Windows Server 2012 R2
	Windows Server 2016	<ul style="list-style-type: none"> • Windows Server 2008 R2 • Windows Server 2012 • Windows Server 2012 R2 • Windows Server 2016
SR-IOV Ethernet	Windows Server 2012	<ul style="list-style-type: none"> • Windows Server 2012

Table 1 - Supported Operating Systems

Virtualization Mode	Supported Host OS	Supported Guest OS
SR-IOV Ethernet	Windows Server 2012 R2	<ul style="list-style-type: none"> Windows Server 2012 Windows Server 2012 R2 Windows 8.1 (64 bit only) Windows client 10 (64 bit only)
	Windows Server 2016	<ul style="list-style-type: none"> Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 Windows Client 8.1 (64 bit only) Windows Client 10 (64 bit only)
SR-IOV InfiniBand KVM	KVM	<ul style="list-style-type: none"> Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2
SR-IOV Ethernet KVM	KVM	<ul style="list-style-type: none"> Windows Server 2012 Windows Server 2012 R2

- a. When installing WinOF v5.22 and above over Windows 7 Client, Microsoft Security Advisory 3033929 is required (can be downloaded from the following location:
<https://technet.microsoft.com/en-us/library/security/3033929?f=255&MSPPErr=2147217396>).
 For further information, please refer to the User Manual.

1.4 Supported Network Adapter Cards

Mellanox WinOF Rev 5.35 supports the following Mellanox network adapter cards:

Table 2 - Supported Network Adapter Cards

NICs	Supported Protocol	Supported Link Speed
ConnectX®-3 Pro	InfiniBand (IB)	SDR, DDR, QDR, FDR10, FDR
	Ethernet	10, 40, 50 and 56 Gb/s
ConnectX®-3	InfiniBand (IB)	SDR, DDR, QDR, FDR10, FDR
	Ethernet	10, 40, 50 and 56 Gb/s



The speeds listed in the table above are according to the speeds supported by the device.

1.4.1 Firmware Versions

Mellanox WinOF Rev 5.35 provides the following firmware for Mellanox NICs:

Table 3 - Firmware Versions

NICs	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3 Pro / ConnectX®-3 Pro EN	2.40.5032	N/A
ConnectX®-3 / ConnectX®-3 EN	2.40.5032	N/A



Firmware version 2.40.5000 requires upgrade to version 2.40.5030 and above, before installing driver version 5.35 and above.

1.5 RoCE Modes Matrix

The following is RoCE modes matrix.

Table 4 - RoCE Modes Matrix

Software Stack / Inbox Distribution	RoCE MAC Based (Layer 2) Supported as of Version	RoCE IP Based (Layer 2) Supported as of Version	RoCE v2 (Layer 3) Supported as of Version
Mellanox WinOF	3.2 (Default)	4.80 (Requires additional configuration)	<ul style="list-style-type: none"> 4.70 (Requires additional configuration) 5.22 (Default)
Inbox Windows Server 2012 / Inbox Windows Server 2012 R2	Supported (Default)	Not supported	Not supported
Inbox Windows Server 2016	Supported	Supported	Supported (Default)



As of WinOF 5.22, RoCE v2 is the default RoCE mode.

Table 5 - RoCE v2 UDP Port Matrix

WinOF Versions	RoCE v2 UDP port
4.70 - 5.00	1021
5.10 and above	4791

For further information, please refer to section “RoCE v2 UDP Port” in the User Manual.

2 Changes and New Features in Rev 5.35



This package version is 5.35.52000. The package contains the following versions of components:

- Bus, eth, IPoIB and mux drivers version is 5.35.12978
- The CIM provider version is 5.35.12978

Table 6 - Changes and New Features in Rev 5.35

Category	Description
Ethernet	Updated driver settings for Virtual Function to receive optimal performance
	Added RSS (Rx Steering Mode) monitoring counters support
	Added counter for link up/down to count the number of times that the link operative state changes to down. See “ <i>Proprietary Mellanox Adapter Diagnostics Counters</i> ” in the User Manual.
Tools	vstat tool - Added interface description for each port in the vstat tool.
	Mlxtool - Added support in the following actions:
	• Driver restart
	• Pulling resource tracker information for VFs
Bug fixes	See “Bug Fixes History” on page 37.

For further information on the new features stated here, please refer to the WinOF User Manual.

2.1 Beta Features

Table 7 - Beta Features

Category	Description
ibdump	ibdump is currently at beta level.
IPoIB	IPv6 support of IPoIB in an SR-IOV guest OS over KVM is at beta level.
	IPoIB teaming support is at beta level and is supported only on native machines (and not in HyperV or SR-IOV).
	VMQ over IPoIB in Windows Server 2016 is at beta level due to Known Issue #946432, see “Known Issues” on page 12.
Virtualization	RDMA over VM when in SR IOV mode is currently at beta level.

2.2 Unsupported Functionality/Features

The following are the unsupported functionalities/features in WinOF:

- ND over WinVerbs provider
- SRP
- IPv6 over IPoIB Team ports
- VMQ over IPoIB team ports
- Configure IPoIB team through PowerShell
- ConnectX®-2 adapter cards
- IBAL performance tools (ib* ibv*)
- WinVerbs
- IBVerbs

3 Known Issues

The following table provides a list of known bugs and limitations in regards to this release of WinOF.

Table 8 - Known Issues (Sheet 1 of 25)

Internal Ref.	Issue
-	<p>Description: 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.</p> <p>WA: Avoid pinning the whole machine memory in those Operating Systems.</p> <p>Keywords: Generic</p>
-	<p>Description: When running applications that use ND or libibmad (such as OpenSM) the system might get to an unstable state when trying to shutdown/restart/hibernate it.</p> <p>WA: Close all applications that use ND or libibmad before performing shutdown/restart/hibernate.</p> <p>Keywords: Generic</p>
322721	<p>Description: Activating NC-SI in WinOF v4.90.10541 may cause driver's loading failure when using an older firmware version than 2.30.8000.</p> <p>WA: Do not enable NC-SI in machines that WinOF v4.90.10541 is installed in.</p> <p>Keywords: Generic</p>
-	<p>Description: 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.</p> <p>WA: N/A</p> <p>Keywords: Generic</p>
-	<p>Description: Running <code>Ntttcp</code> without the “-a X” flag ($X > 1$) in a NIC configured with 10GbE, may cause low bandwidth in TCP single stream.</p> <p>WA: Run <code>Ntttcp</code> with “-a 8” for best performance</p> <p>Keywords: Generic</p>
-	<p>Description: Active links disappear after changing the cable connectivity from Ethernet to InfiniBand or vice versa.</p> <p>WA: Disable and enable the <code>mlx4_bus</code> interface from the Device Manager.</p> <p>Keywords: Generic</p>

Table 8 - Known Issues (Sheet 2 of 25)

Internal Ref.	Issue
-	<p>Description: 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 ibstat). Mlxburn/flint return 0xffff 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>WA: Use the GUID value returned by the fabric/driver utilities (not 0xffff).</p> <p>Keywords: Generic</p>
-	<p>Description: The UI displays the network interface's state incorrectly even after disabling and re-enabling the "disconnected" network interfaces.</p> <p>WA: To see the correct state and link speed, perform one of the following:</p> <ul style="list-style-type: none"> • Run Get-netadapter Powershell cmdlet <p>or</p> <ul style="list-style-type: none"> • Right click on that interface from "Network Connections" and click on status <p>Keywords: Generic</p>
-	<p>Description: WoL is not supported on Windows Server 2008 R2.</p> <p>WA: N/A</p> <p>Keywords: Generic</p>
-	<p>Description: iSCSI boot over Windows Server 2008 R2 is not supported.</p> <p>WA: N/A</p> <p>Keywords: Generic</p>
410269	<p>Description: Clearing the Mellanox device counters through Perfmon does not always work.</p> <p>WA: Restart the driver</p> <p>Keywords: Generic</p>
401792	<p>Description: When running in SR-IOV mode and Hyper-V, the same driver version should be used on both the Hypervisor and the virtual machine.</p> <p>WA: N/A</p> <p>Keywords: Generic</p>
-	<p>Description: Working with user space RDMA resources (i.e. ND) in parallel with hardware reset may cause unexpected behavior.</p> <p>WA: N/A</p> <p>Keywords: Generic</p>

Table 8 - Known Issues (Sheet 3 of 25)

Internal Ref.	Issue
491668	Description: In Windows Server 2008 R2 and Windows 7 Client, ConnectX-3 Pro is identified as ConnectX-3. This does not affect system behavior.
	WA:
	Keywords: Generic
-	Description: 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.
	WA: N/A
	Keywords: Generic
-	Description: The drivers' functionality is limited up to 128 cores.
	WA: N/A
	Keywords: Generic
-	Description: According to the hardware architecture, Mellanox NIC devices are presented in the registry as virtual devices and not as physical devices.
	WA: Any customer programs or scripts that look into the NIC registry values should not assume it is a physical device.
	Keywords: Generic
464449	Description: IPv6 traffic between Hyper-V hosts over IPoIB v-Switch may experience traffic loss.
	WA: N/A
	Keywords: InfiniBand
491546	Description: Creating a virtual IPoIB port with non-default PKey isn't supported in the WinOF upgrade. Doing so will indicate that the adapter is unplugged after the upgrade ends.
	WA: Remove the virtual port before the upgrade and re-add it after the upgrade.
	Keywords: InfiniBand
-	Description: InfiniBand application that was compiled with an SDK version earlier than WinOF v4.90 is not binary compatible.
	WA: Recompile InfiniBand application with WinOF v4.90 and above. ND application is backward compatible and older applications over ND do not have to be recompiled.
	Keywords: InfiniBand

Table 8 - Known Issues (Sheet 4 of 25)

Internal Ref.	Issue
186806	<p>Description: WinOF 4.40 and above IPoIB driver is not IPv6 compatible with earlier driver versions, including Window 8 Inbox driver. If WinOF 4.50 IPoIB node receives an icmpv6 message from the legacy IPoIB node, the following event will appear in the event log: “IPoIB driver detected a possible incompatibility with Windows 8 inbox IPv6 support due to which there is no IPv6 connectivity”.</p> <p>WA: 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. We recommend upgrading all hosts to the new driver version from http://mellanox.com or use Windows Server 8 compatibility mode</p> <p>Keywords: InfiniBand</p>
-	<p>Description: Without separate ports for each stream, WinSock multiplexes every packet to every subscriber socket and then filters it out.</p> <p>WA: Use different UDP ports to get higher performance when using multicast packets.</p> <p>Keywords: InfiniBand</p>
-	<p>Description: A virtual IPoIB interface, created by the part_man utility, reports an Active state when the physical link is in the Initializing state and OpenSM is not running in the subnet.</p> <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: The “Packets Received Discarded” and “Packets Received Errors” counter may display wrong results.</p> <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: Connection failure on ND tests while machine A have IBAL provider and machine B have MLX4ND provider.</p> <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: Hibernate and Sleep are not functional when user-space is using its resources.</p> <p>WA: N/A</p> <p>Keywords: InfiniBand</p>

Table 8 - Known Issues (Sheet 5 of 25)

Internal Ref.	Issue
-	<p>Description: IPoIB does not support:</p> <ul style="list-style-type: none"> • MAC address change • QoS (packet priority) • Connected Mode <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: 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.</p> <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: 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.</p> <p>WA: To run OpenSM as a service, assuming the package was installed in the default path, use: sc create opensm binPath= "c:\Program Files\Mellanox\MLNX_VPIB\Tools\opensm.exe" To start the service run: sc start opensm</p> <p>Keywords: InfiniBand</p>
-	<p>Description: Tools issues:</p> <ul style="list-style-type: none"> • ibportstate does not work on RoCE ports • ibdiagpath may crash on Hyper-V machines <p>WA: N/A</p> <p>Keywords: InfiniBand</p>
-	<p>Description: If an application which uses InfiniBand runs while the driver is being restarted, a bluescreen or an NMI may occur.</p> <p>WA: Stop all InfiniBand applications including OpenSM upon driver restart.</p> <p>Keywords: InfiniBand</p>
330284	<p>Description: If OpenSM is up during driver restart on the same machine, it might stop working, and an error message that reads “Another OpenSM is running” will pop up.</p> <p>WA: To resume operation, stop OpenSM and restart the driver while OpenSM is down.</p> <p>Keywords: InfiniBand</p>

Table 8 - Known Issues (Sheet 6 of 25)

Internal Ref.	Issue
-	Description: No communication between the physical interface and a VM that uses vSwitch created over virtual IPoIB (Pkey), and vice versa.
	WA: N/A
	Keywords: InfiniBand
90224	Description: Burning a firmware image with a "mtu_cap" value other than the default causes the driver load to fail.
	WA: Do not change the "mtu_cap" value
	Keywords: InfiniBand
383460	Description: Calling <code>ib_query_cq()</code> on a CQ which was created with 1 entry indicates that there are 0 entries in the CQ.
	WA: N/A
	Keywords: InfiniBand
439805	Description: IPoIB ports report the configured MTU instead of the effective MTU. These MTUs differ when the corresponding InfiniBand partition is configured with a smaller MTU than the port's configured MTU. In such case, there may be communication failures and/or CPU work overhead on the communicating nodes (even if the peer node is configured correctly).
	WA: Make sure the configured adapter MTU (the "Jumbo packet" property in the adapter's advanced settings) is consistent with the corresponding InfiniBand partition MTU. Specifically, the default IPoIB adapter MTU is 4KB and the default OpenSM partition MTU is 2KB. Thus, the default MTU of adapters and OpenSMs is inconsistent. Therefore, the MTU of either all adapters or all OpenSMs (which may be on network switches) on the InfiniBand subnet needs to be changed. Note that the adapter's MTU can be configured through PowerShell.
	Keywords: InfiniBand
-	Description: 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.
	WA: N/A
	Keywords: InfiniBand
498130	Description: DHCP messages that IPoIB send are not fully spec compliant. The spec specifies that the 'chaddr' field must be zeroed, but WinOF IPoIB does not guarantee it.
	WA: N/A
	Keywords: InfiniBand

Table 8 - Known Issues (Sheet 7 of 25)

Internal Ref.	Issue
-	<p>Description: Disabling the “Priority & VLAN tag” in the UI which VLANID is configured, may result in sending packets with the configured VLANID.</p> <p>WA: Remove the VLANID before disabling the “Priority & VLAN tag”.</p> <p>Keywords: Ethernet</p>
206528/ 206945	<p>Description: WakeOnMagicPacket registry key is not added to the registry although WoL is supported by the driver and by the NIC.</p> <p>WA: N/A</p> <p>Keywords: Ethernet</p>
189704	<p>Description: When the ports of the device are configured as Ethernet only, ibstat/vstat may display wrong information.</p> <p>WA: N/A</p> <p>Keywords: Ethernet</p>
-	<p>Description: High multicast drop rate on multicast storming.</p> <p>WA: Use “Multicast traffic” tuning option under the performance tab. For further information, please refer to section “Tunable Performance Parameters” in the User Manual.</p> <p>Keywords: Ethernet</p>
-	<p>Description: When there is a stress in TCP connection establishments, some of those connections may fail.</p> <p>WA: Increase the Ring queue sizes:</p> <ul style="list-style-type: none"> • ReceiveBuffers - controls the receive ring size • TransmitBuffers - controls the transmit ring size <p>Keywords: Ethernet</p>
-	<p>Description: 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:</p> <ul style="list-style-type: none"> • It has a traffic classification condition of type NDIS_QOS_CONDITION_DEFAULT. • It is the first traffic classification defined in the array of NDIS_QOS_CLASSIFICATION_ELEMENT structures. <p>WA: N/A</p> <p>Keywords: Ethernet</p>

Table 8 - Known Issues (Sheet 8 of 25)

Internal Ref.	Issue
416248/ 409093	Description: Disabling the Ethernet adapter with more than 25 VLANs configured over Windows Server 2008 R2 \ Windows 7 Client, may result with a non-responding server.
	WA: Reduce the number of configured VLANs prior to disabling the Ethernet adapter.
	Keywords: Ethernet
410377	Description: Virtual Ethernet Interfaces created by vea_man are not tuned by the automatic performance tuning script.
	WA: For optimal performance need to follow the performance tuning guide and apply relevant changes to the VEA interface.
	Keywords: Ethernet
-	Description: In Windows Server 2008 R2, devices created by the Teaming driver do not show the correct OEM branding in the Device Manager.
	WA: N/A
	Keywords: Ethernet
326885	Description: Wake-on-Lan (WoL) cannot be disabled on NICs which supports it.
	WA: N/A
	Keywords: Ethernet
-	Description: Running Quality of Service (QoS) commands without the parameter “-PolicyStore ActiveStore” may cause machines to load without Quality of Service policy.
	WA: Store the QoS policy in the ActiveStore
	Keywords: Quality of Service
-	Description: 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
	WA: N/A
	Keywords: RoCE NVGRE

Table 8 - Known Issues (Sheet 9 of 25)

Internal Ref.	Issue
327531	Description: In machines with heterogeneous NICs, a NIC which supports RoCE v2, and a NIC which does not support RoCE v2 the following issues might raise: <ul style="list-style-type: none"> • ConnectX®-3 Pro is loaded with the transport type RoCE v2 • ConnectX®-3 Pro is loaded with the transport type RoCE
	WA: N/A
	Keywords: RoCE NVGRE
-	Description: When NVGRE off-load is enabled, the GRE traffic cannot be accepted as a regular L2 traffic and requires special L2_TUNNELING steering rules. In such case the GRE packets are dropped or directed to promiscuous queue.
	WA: N/A
	Keywords: RoCE NVGRE
193619	Description: Using different versions of RoCE in your cluster is not supported.
	WA: Use the same RoCE version in all the cluster in the Ethernet ports.
	Keywords: RoCE NVGRE
-	Description: RDMA Activity counters do not count during NetworkDirect RoCE traffic.
	WA: N/A
	Keywords: RoCE NVGRE
-	Description: GRE traffic steering by inner MAC and by outer MAC simultaneously is currently not supported.
	WA: Configure steering or by inner MAC, or by outer MAC.
	Keywords: RoCE NVGRE
-	Description: 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.
	WA: N/A
	Keywords: RoCE NVGRE
-	Description: Set the bus driver registry keyAcceptGREbyOuterMAC_P1/2 per port to accept GRE traffic by outer MAC and to duplicate L2 steering rule to L2_TUNNELING 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 0:AcceptGREbyOuterMAC_P1/2
	WA: N/A
	Keywords: RoCE NVGRE

Table 8 - Known Issues (Sheet 10 of 25)

Internal Ref.	Issue
206696	Description: When using WinOF 4.40 or above, low throughput will be seen on 40GbE adapters when QoS is enabled.
	WA: Disable QoS when it is not in use. Open a PowerShell prompt. Run: Disable-NetAdapterQos -name <Interface Name> where <Interface Name> is e.g. "Ethernet 1"
	Keywords: Performance
-	Description: perf_tuning is supported only when one of the two NUMA nodes are in use.
	WA: N/A
	Keywords: Performance
-	Description: Running performance benchmarks for a short period of time (< 1 sec) may provide bad latency in IPoIB and Ethernet.
	WA: Set "Rx Interrupt Moderation Profile" and "Tx Interrupt Moderation Profile", to "Low Latency" to avoid bad latency. Note: This may increase CPU utilization.
	Keywords: Performance
-	Description: The driver uses optimal interrupt moderation values for 10 GbE SR-IOV VF scenario. For other scenarios, the optimal values yet to be found.
	WA: N/A
	Keywords: Performance
443137/ 439897	Description: While running in a Virtual Machine (working with VF) or in a Native Machine, performance counters, if read directly or by using any tool (as Windows Task Manager), may show that no packet has been sent/received. This happens because the driver periodically examines the actual performance counters and caches the results. If the sample rate is too high, the counter values will remain the same.
	WA: Reduce the counters sample rate.
	Keywords: Performance
-	Description: 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.
	WA: N/A
	Keywords: Hyper-V

Table 8 - Known Issues (Sheet 11 of 25)

Internal Ref.	Issue
-	<p>Description: 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> <p>WA: Reattach the vSwitch to sync with the value set in the ETH\IPoIB device.</p> <p>Keywords: Hyper-V</p>
-	<p>Description: 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>WA: N/A</p> <p>Keywords: Hyper-V</p>
-	<p>Description: 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>WA: N/A</p> <p>Keywords: Hyper-V</p>
-	<p>Description: Hyper-V is at low bandwidth on LBFO vSwitch, Windows Server 2012.</p> <p>WA: N/A</p> <p>Keywords: Hyper-V</p>
-	<p>Description: 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> <p>WA: N/A</p> <p>Keywords: Hyper-V</p>
-	<p>Description: 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.</p> <p>WA: N/A</p> <p>Keywords: Hyper-V</p>

Table 8 - Known Issues (Sheet 12 of 25)

Internal Ref.	Issue
-	<p>Description: 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.</p> <p>WA: N/A</p> <p>Keywords: Hyper-V</p>
-	<p>Description: The IPoIB non-VMQ mode is supported only when the VMQ is enabled according to the registry values.</p> <p>WA: To use the non-VMQ mode for a VM, change its settings as follow:</p> <ul style="list-style-type: none"> • Press “Settings” on the VM • Go to Network Adapter -> Hardware Acceleration • Un-check the “Enable virtual machine queue” <p>Keywords: Hyper-V</p>
-	<p>Description: Working with Windows guest OS over non-Windows SR-IOV hypervisor may result with higher latency compared to Windows Hypervisor.</p> <p>WA: N/A</p> <p>Keywords: SR-IOV</p>
-	<p>Description: When working with OEM adapter cards in SR-IOV mode, the VF is identified as a Mellanox adapter. There is no other effect on the behavior of the VF.</p> <p>WA: Verify that the FW INI contains the parameter vf_subsystem_id with the OEM encoding in the [HCA] section.</p> <p>Keywords: SR-IOV</p>
-	<p>Description: 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.”</p> <p>WA: Start the OpenSM and restart the driver.</p> <p>Keywords: SR-IOV</p>
418268	<p>Description: For InfiniBand SR-IOV guest, OpenSM Assigned GUIDs are not supported and may cause unexpected behavior.</p> <p>WA: Work only with Administrator assigned GUIDs.</p> <p>Keywords: SR-IOV</p>

Table 8 - Known Issues (Sheet 13 of 25)

Internal Ref.	Issue
408734	Description: For InfiniBand SR-IOV guest, ND traffic does not work on an SR-IOV adapter when a Para-Virtualization adapter configured on the same virtual machine and the same subnet as IPoIB.
	WA: Disable the Para-Virtualization adapter.
	Keywords: SR-IOV
-	Description: In SR-IOV mode over Hyper-V, all ports are in pure Ethernet mode and RDMA is not supported on either port.
	WA: N/A
	Keywords: SR-IOV
342421	Description: 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.
	WA: Unbond all Vswitches from the NIC's interfaces before enabling SR-IOV in the BIOS.
	Keywords: SR-IOV
427038	Description: 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.
	WA: N/A
	Keywords: SR-IOV
-	Description: Device name in the Device Manager does not show the correct OEM branding for SR-IOV Virtual Function devices without the updated firmware.
	WA: N/A
	Keywords: SR-IOV
-	Description: Upgrading the driver while the UI is opened with the “ConnectX NIC device” may cause the installation process to never end.
	WA: Close the UI before driver upgrade.
	Keywords: Installation / Upgrade
207497	Description: Rebooting the machine while uninstalling WinOF may result in installation failure.
	WA: Delete Mellanox components from HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\DIFxApp\Components. The Mellanox components are mlx4eth63, ipoib6x and mlx4_bus
	Keywords: Installation / Upgrade

Table 8 - Known Issues (Sheet 14 of 25)

Internal Ref.	Issue
137859	Description: 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)".
	WA: Uninstall the current version and install the older one.
	Keywords: Installation / Upgrade
198537	Description: Downgrade is not supported.
	WA: N/A
	Keywords: Installation / Upgrade
-	Description: Configuration is not restored when replacing a ConnectX®-3 NIC with a ConnectX®-3 Pro NIC located on the same PCI slot
	WA: Clean the old network adapter configuration prior to upgrade.
	Keywords: Installation / Upgrade
-	Description: Configuration can be restored only in Windows Server 2012 and above.
	WA: N/A
	Keywords: Installation / Upgrade
-	Description: IPv6 configuration restore is not supported.
	WA: N/A
	Keywords: Installation / Upgrade

Table 8 - Known Issues (Sheet 15 of 25)

Internal Ref.	Issue
-	<p>Description: Upon upgrade, the following Registry Key values will be over-written 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>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: Upon upgrade the following Ethernet Registry Keys will be deleted:</p> <ul style="list-style-type: none"> • SendCompletionMethod • UseRSSForRawIP • UseRSSForUDP <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: Upon upgrade the SendCompletionMethod IPoIB Registry Key value will be modified as follow: SendCompletionMethod = 0</p> <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: Upon upgrade the following IPoIB Registry Keys will be deleted:</p> <ul style="list-style-type: none"> • UseRSSForRawIP • UseRSSForUDP <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>

Table 8 - Known Issues (Sheet 16 of 25)

Internal Ref.	Issue
383584	Description: 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.
	WA: 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.
	Keywords: Installation / Upgrade
-	Description: Running a downgrade in silent mode is not supported. Upon downgrade the return code will always be 0.
	WA: N/A
	Keywords: Installation / Upgrade
-	Description: Uninstalling the driver after upgrade won't remove the directory %ProgramFiles%\Mellanox
	WA: N/A
	Keywords: Installation / Upgrade
-	Description: Uninstalling the driver when multiple of VLANs are configured never ends.
	WA: Remove the VLANs before uninstallation.
	Keywords: Installation / Upgrade
-	Description: The installation process does not close any applications running in the background, and may cause a BSOD as a result of a stuck cmd.
	WA: It is recommended to close all running applications prior to upgrading the driver.
	Keywords: Installation / Upgrade
-	Description: 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>\AppData\Local\Temp\DMI151A.tmp.log.xml C:\Program Files\Mellanox\MLNX_VPI\ETH\mlx4eth63.inf
	WA: Reboot the machine and reinstall
	Keywords: Installation / Upgrade

Table 8 - Known Issues (Sheet 17 of 25)

Internal Ref.	Issue
403352	Description: 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”
	WA: Kill the WMIPrvSE.exe tasks in the task manager and reinstall.
	Keywords: Installation / Upgrade
-	Description: 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. 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
	WA: Run the following script on the server to optimizes loading PowerShell DLLs: \$Env:PATH = [Runtime.InteropServices.RuntimeEnvironment]::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 } Keywords: Installation / Upgrade
	Keywords: Installation / Upgrade
417380/ 415257	Description: On ConnectX®-3 cards only, when upgrading from Windows Server 2012 R2 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.
	WA: Reconfigure the RoCE Mode setting manually.
	Keywords: Installation / Upgrade
427676	Description: Uninstall from the Device Manager is currently unsupported
	WA: Driver uninstall can be completed from the Programs and Features window as explained in the “Uninstalling Mellanox WinOF Driver” section in the User Manual.
	Keywords: Installation / Upgrade

Table 8 - Known Issues (Sheet 18 of 25)

Internal Ref.	Issue
-	<p>Description: 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.</p> <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: 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.</p> <p>WA: Upgrade the firmware manually.</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: Driver installation requires deletion of the mlx4_bus.sys file in Windows Server 2008 R2 when using the PXE package.</p> <p>WA: Delete the mlx4_bus.sys file and reboot the machine to install the driver.</p> <p>Keywords: Installation / Upgrade</p>
492398	<p>Description: If there are disabled network interfaces or a disabled Mellanox bus driver, they will be enabled after the WinOF upgrade.</p> <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: The upgrade from WinOF v4.58 to WinOF v4.90 saves only the configuration of the Ethernet interface, and does not save the configuration of the IPoIB interface.</p> <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
647102	<p>Description: Using a non-Mellanox HWID will result in losing all NIC configurations when upgrading from a version older than v4.80</p> <p>WA: N/A</p> <p>Keywords: Installation / Upgrade</p>
-	<p>Description: ibdump may encounter packet drops upon a burst of more than 4096 (or 2^{max-burst}) packets.</p> <p>WA: N/A</p> <p>Keywords: Utilities</p>

Table 8 - Known Issues (Sheet 19 of 25)

Internal Ref.	Issue
-	Description: Packets loss is not reported by ibdump.
	WA: N/A
	Keywords: Utilities
-	Description: 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.
	WA: Verify Ethernet flow control is enable to ensure a lossless link
	Keywords: Utilities
-	Description: Pcap file issues for RoCE IP Based: <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'.
	WA: N/A
	Keywords: Utilities
-	Description: Sniffing over IB ports is currently not supported
	WA: N/A
	Keywords: Utilities
199079	Description: When the tracer tool traces many events, it may consume a large amount of memory (up to several GB RAM).
	WA: Use the button to set maximum number of the displayed lines at a maximum buffer size.
	Keywords: Utilities
367772	Description: ibdump works only on Ethernet links.
	WA: N/A
	Keywords: Utilities
-	Description: Running Microsoft CIM cmdlets operations and their derived classes on classes MLNX_NetAdapterSettingData and MLNX_NetAdapterRoceSettingData is not supported. Calling those commands may cause the debugger, if connected to the machine, to assert.
	WA: Use DriverCoreSettings instead.
	Keywords: CIM/WMI
408230	Description: For PCI Gen3, PcieLinkSpeed is reported as “Unknown” when running Get-NetAdapterHardwareInfo Powershell cmdlet
	WA: N/A
	Keywords: CIM/WMI

Table 8 - Known Issues (Sheet 20 of 25)

Internal Ref.	Issue
-	Description: WMI does not work due to lack of permissions.
	WA: Change the execution policy. Run: Set-ExecutionPolicy AllSigned
	Keywords: CIM/WMI
433986	Description: The information that is printed in the cmdlets get-netadaptersriov and Get-MlnxPCIDeviceSetting is inconsistent.
	WA: Use only the cmdlet Get-MlnxPCIDeviceSetting.
	Keywords: CIM/WMI
432674	Description: There is an interoperability problem between NDK and ND in RDMA operations during the write and read operations. However, the send operation resumes working. This happens since ND uses remote token in a network order (big endian) while NDK uses remote token in CPU order (little endian). Therefore, an inconsistency is caused between ND and NDK in RDMA operations.
	WA: An ND application that works with NDK using RDMA operations must handle this issue by changing the remote token to the appropriate byte order before sending it to NDK.
	Keywords: ND
-	Description: When working with the default NDv1 and NDv2 providers, the following error message might be displayed: 0xC0000120 NT_STATUS_CANCELLED This error does not affect any functionality and can be safely ignored.
	WA: N/A
	Keywords: ND
-	Description: Changing the default ND providers may cause random errors, such as: 0xC0000238 (NT_STATUS_ADDRESS_ALREADY_ASSOCIATED) on Connect() or with 0xC0000236 (NT_STATUS_CONNECTION_REFUSED) on Accept(). These errors can be safely ignored.
	WA: N/A
	Keywords: ND
759661	Description: UDP latency increases from 13.5usec to 15.9 when ECN is enabled.
	WA: N/A
	Keywords: Performance

Table 8 - Known Issues (Sheet 21 of 25)

Internal Ref.	Issue
464449/ 464452/ 464453/ 464454	Description: IPoIB IPv6 over KVM/ESXi is not supported
	WA: N/A
	Keywords: Virtualization
798900	Description: CIM command does not function properly in VM in SR-IOV mode.
	WA: N/A
	Keywords: CIM command, VM, SR-IOV
798219	Description: WinOF contains hardware-specific performance optimization in the form of reduced Memory Translation Table (MTT) footprint for Network Direct applications' Memory Region (MR) buffers (created through IND-Adapter::RegisterMemory() method). This optimization is not implemented for MR buffers that are larger than 2GB.
	WA: N/A
	Keywords: RDMA
800648	Description: On Hyper-V machines with SR-IOV VSwitch, the amount of resources available to the host for RDMA purposes is limited.
	WA: Detach the SR-IOV VSwitch from the Mellanox NIC
	Keywords: SR-IOV
798033/ 817741	Description: Virtual Ethernet Adapters (VEA) disabled by default since SR-IOV is enabled by default. When upgrading to WinOF v5.22, Virtual Ethernet Adapters will become nonfunctional.
	WA: Before upgrading from an earlier version than WinOF v5.22, make sure to disable SR-IOV. For further information on the registry key, please refer to the “SR-IOV Registry Keys” section in WinOF User Manual.
	Keywords: SR-IOV, Ethernet, VEA, Virtual Ethernet Adapter
829865	Description: Driver fails to load on a server with more than 4TB Ram.
	WA: 1. Set the following registry key: Reg path: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\mlx4_bus\Parameters Reg key name: LogNumMttOverride Reg key type: DWORD Reg key value: 1 b. Restart the device.
	Keywords: mlx4_bus, Load failure

Table 8 - Known Issues (Sheet 22 of 25)

Internal Ref.	Issue
843240	Description: Guest RDMA is supported only in a VM with a single VF.
	WA: N/A
	Keywords: RDMA, NDK V2, ND
-	Description: The driver fails to load on guest machine with Windows 10 and Inbox driver 4.91 when in SR-IOV mode.
	WA: Since SR-IOV mode in v5.25 driver does not support Windows 10 old Inbox driver (4.91), we recommend you update your driver with the latest version of Windows 10.
	Keywords: SR-IOV, Windows 10
-	Description: LBFO over IPoIB is not supported.
	WA: N/A
	Keywords: LBFO, IPoIB
848675	Description: NDK on the Guest Virtual Machine is not functional after restarting the bus driver in the host.
	WA: Restart the bus driver in the Guest Virtual Machine
	Keywords: LBFO, IPoIB
850804	Description: When RDMA is enabled on a Virtual Machine in SR-IOV mode, vstat will always show that port state is done even though the port is up.
	WA: N/A
	Keywords: RDMA, VM, SR-IOV, vstat
884986	Description: The ucrtbse.dll file is missing in operating systems older than Windows Server 2016. This leads to failure of the installation. The UI and installation log will show messages for all missing updates
	<ul style="list-style-type: none"> • UI Message Example: "missing windows updates: KB2999226" • Log File Message Examples: "CheckWinUpdate: Failed to locate windows update 2999226" "CheckWinUpdate: please refer to: https://support.microsoft.com/en-us/kb/2999226".
	WA: Install Windows update kb2999226. The update files can be downloaded from https://support.microsoft.com/en-us/kb/2999226 .
	Keywords: Windows Server 2008 R2, Windows Server 2012, Windows Server 2012R2, Teaming, vlan_config, api-ms-win-core-timezone-l1

Table 8 - Known Issues (Sheet 23 of 25)

Internal Ref.	Issue
805148	Description: The Memory Translation Table (MTT) optimization feature is disabled for memory regions larger than 2GB.
	WA: N/A
	Keywords: Memory, MTT
909038	Description: Windows 2016 mistakenly recognizes ConnectX-3 Pro card as a USB interface. This may disconnect the network and cause failure of critical network services.
	WA: N/A
	Keywords: ConnectX-3 Pro, network services failure
909040	Description: When configuring ConnectX-3 Pro single port via the driver, the RDMA connection fails.
	WA: N/A
	Keywords: ConnectX-3 Pro, single port, RDMA
909274	Description: RoCE fallback mode is not working when SR-IOV is enabled.
	WA: N/A
	Keywords: RoCE, fallback mode, SR-IOV
648334	Description: The IPoIB driver may stall upon disabling, due to a non-zero reference count in TCP/IP stack.
	WA: N/A
	Keywords: IPoIB, disable
909595	Description: Errors that may occur during initial driver start in the DriverEntry function do not appear in the event log. These errors will be indicated by a yellow exclamation mark on the device, with code 37.
	WA: Parsing the etl file provides the required information
	Keywords: mlx4_bus
935344	Description: Several error messages from the mlx4_bus may appear in the VM Event Log. These messages (104, 122, 144) do not indicate any error, and can be safely ignored.
	WA: N/A
	Keywords: Event log
951028	Description: Ethernet interface is started, in spite of InfiniBand configuration of the port. This occurs due to an enabled BMC or an outdated firmware version.
	WA: Disable BMC, or update the installed firmware version.
	Keywords: port type, VPI, InfiniBand

Table 8 - Known Issues (Sheet 24 of 25)

Internal Ref.	Issue
961699	Description: On hypervisor, when one Ethernet port is bind to vmswitch in which SR-IOV is enabled, Network Direct applications do not work on the second port.
	WA: N/A
	Keywords: hypervisor, vmswitch, SR-IOV
948671	Description: Vlan_config tool allows to modify the vlan name to a name which is invalid by the vlan_config addvlan command.
	WA: N/A
	Keywords: Vlan_config, vlan name, addvlan
944624	Description: A VF can initialize only one Comm Channel: either MS Comm Channel or MLNX Comm Channel.
	WA: N/A
	Keywords: VF, Comm Channel
936651	Description: When removing a PKey that is a part of an IPoIB team interface, the team is disabled and cannot be deleted.
	WA: Remove the IPoIB team before deleting the PKey.
	Keywords: IPoIB Pkeys, Team
801288	Description: On Hyper-V machines with SR-IOV VSwitch, the amount of resources available to the host for RDMA purposes is limited.
	WA: Detach the SR-IOV VSwitch from the Mellanox NIC.
	Keywords: Hyper-V, SR-IOV VSwitch, RDMA
839712	Description: The tools' folder is not added to the path environment variable when installing the driver on a VM.
	WA: After installing the driver, perform one of the following workarounds: 1. Go to the Control Panel --> System and Security --> System Click on "Advanced System Settings" (on the right panel) Select "Advanced" and click on the "Environment Variables" button Click on the OK button in the "Environment Variables" window. 2. Log off or reboot the driver
	Keywords: tools' folder, VM
798245	Description: ConnectX-3/ConnectX-3 Pro bus driver counters and Ethernet driver counters are not available in Windows 7 Client/Windows 2008R2.
	WA: N/A
	Keywords: ConnectX-3, ConnectX-3 Pro, performance counters, Windows 7 Client, Windows 2008R2

Table 8 - Known Issues (Sheet 25 of 25)

Internal Ref.	Issue
964757	Description: On servers where Hyper-v role is installed and SR-IOV is enabled, killing the ibdump process will cause a resource leak.
	WA: Use CTRL-C to stop the ibdump process, instead of CTRL-Break.
	Keywords: ibdump, resource leak
972033	Description: Changing RoCE mode in a virtual machine is not supported.
	WA: N/A
	Keywords: RoCE, Virtual Machine

4 Bug Fixes History

Table 9 lists the bugs fixed in this release.

Table 9 - Fixed Bugs List (Sheet 1 of 17)

Internal Ref.	Issue
995925	Description: Fixed an issue that occurred rarely when working with LSO - a fragmented packet (with more than 41 fragments) could lead to duplicated headers.
	Keywords: LSO, fragmented packets, duplicated headers
	Discovered in Release: 5.30
	Fixed in Release: 5.35.TBD (FUR 2)
991194	Description: Fixed an issue that caused low bandwidth when using Software vRSS.
	Keywords: Performance, low bandwidth, Software vRSS
	Discovered in Release: 5.30
	Fixed in Release: 5.35.TBD (FUR 2)
966761	Description: Fixed an issue that led to non-optimal Out of box performance for virtual function.
	Keywords: Performance, OOB, SRIOV, virtual function
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12970 (FUR 1)
964639	Description: Fixed an issue which caused a firmware upgrade failure (error code 9) during installation, when RoCE was disabled.
	Keywords: RoCE, firmware upgrade, installation
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12970 (FUR 1)
957390	Description: Fixed an issue where Miniport reset could lead to a driver hang when occurred during driver disabling, or to a system crash when occurred during driver shutdown.
	Keywords: Miniport reset, driver disabling, shutdown
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965
954467	Description: Fixed an issue where the link speed of the IPoIB adapter was the actual speed and not the official speed (i.e. 54.3GB/s instead of 56 GB/s).
	Keywords: IPoIB adapter, link speed
	Discovered in Release: 5.25
	Fixed in Release: 5.35.12965

Table 9 - Fixed Bugs List (Sheet 2 of 17)

Internal Ref.	Issue
936607	Description: Fixed an issue where firmware burning failed on servers with Connectx-3 and Connectx-4 devices.
	Keywords: firmware burning, Connectx-3, Connectx-4
	Discovered in Release: 5.22
	Fixed in Release: 5.35.12965
943258	Description: Fixed an issue where Mellanox counters in Perfmon did not work over HP devices.
	Keywords: Mellanox counters, Perfmon, HP devices
	Discovered in Release: 5.22
	Fixed in Release: 5.35.12965
935523	Description: Fixed an issue where link load of ports connected to virtual machines took more than 10 seconds. the issue occurred on a hyper-v VMQ setup with several virtual machines, and after running massive traffic on the virtual machines.
	Keywords: link load, virtual machines, hyper-v VMQ
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965
940166	Description: Fixed an issue where in a HyperV machine with VMs running, the network interface required a restart after returning from a Low Power State (sleep/hibernate).
	Keywords: Hyper-v, VMQ, port restart duration
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965
892647	Description: Fixed an issue that caused the installation process to hang while checking if the RDSH service is installed.
	Keywords: installation, hang, RDSH
	Discovered in Release: 5.22
	Fixed in Release: 5.35.12965
936813	Description: Fixed a driver crash that occurred when the VPORT-ID on the TX packet was invalid.
	Keywords: crash, VPORT-ID, TX packet
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965

Table 9 - Fixed Bugs List (Sheet 3 of 17)

Internal Ref.	Issue
931155	Description: Updated Link Speed reporting when working with IPoIB and booting using WinPE. This issue caused the link to initialize with a 1Gb/s speed instead of the actual speed.
	Keywords: Link Speed, IPoIB, WinPE
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965
931589	Description: Fixed a rare error that caused a freeze in the error flow during the driver's startup.
	Keywords: mlx4_bus, freeze, startup
	Discovered in Release: 5.30
	Fixed in Release: 5.35.12965
929057	Description: Fixed an SR-IOV team failure caused by an unsuccessful adapter parameters update.
	Keywords: adapter parameters, SR-IOV, teaming
	Discovered in Release: 5.20
	Fixed in Release: 5.35.12965
928030	Description: Fixed an issue which caused a crash in the driver properties dialog, in case more than 8 teaming ports were defined.
	Keywords: crash, driver properties dialog, teaming ports
	Discovered in Release: 5.20
	Fixed in Release: 5.35.12965
897752	Description: Fixed an issue which reported a false error for successful netsh tcp settings via performance tuning
	Keywords: netsh tcp, performance tuning
	Discovered in Release: 5.22
	Fixed in Release: 5.30
915558	Description: Fixed a crash on the VF initialization error-flow.
	Keywords: VF initialization
	Discovered in Release: 5.25
	Fixed in Release: 5.30

Table 9 - Fixed Bugs List (Sheet 4 of 17)

Internal Ref.	Issue
909392	Description: Deactivated the RDMA statistics counters query for vPorts for which RDMA is not enabled.
	Keywords: RDMA, vPorts
	Discovered in Release: 5.25
	Fixed in Release: 5.30
---	Description: Removed unsupported OS from an installer condition message for the Windows 2012R2 package.
	Keywords: Installer, Windows 2012R2
	Discovered in Release: 5.25
	Fixed in Release: 5.30
852485	Description: Fixed an issue that caused occasional failures of the execution of <code>OID_QOS_OFFLOAD_CURRENT_CAPABILITIES</code> on Windows 2016.
	Keywords: QOS
	Discovered in Release: 5.20
	Fixed in Release: 5.30
888501	Description: Fixed an issue that caused traffic loss following an upgrade of Windows 2016 VM to 5.20 driver version.
	Keywords: Upgrade, traffic loss
	Discovered in Release: 5.20
	Fixed in Release: 5.30
902856	Description: Fixed the issue which caused the failure of the powershell command <code>Get_MLNXNetAdapterSettings</code> and the command <code>Get_MLNXNetAdapterFlowControlSettings</code> on servers with Connectx3/Pro and ConnectX4/LX devices.
	Keywords: CIM, powershell, <code>Get_MLNXNetAdapterSettings</code> , <code>Get_MLNXNetAdapterFlowControlSettings</code> ,
	Discovered in Release: 5.20
	Fixed in Release: 5.30
842752	Description: Fixed a crash on the initial stage of the driver's startup.
	Keywords: <code>mlx4_bus</code> , crash on startup
	Discovered in Release: 5.25
	Fixed in Release: 5.30

Table 9 - Fixed Bugs List (Sheet 5 of 17)

Internal Ref.	Issue
844372	Description: Fixed an issue where the SrioPorts settings were not deleted in the process of the driver's uninstallation.
	Keywords: Uninstallation, SrioPorts settings
	Discovered in Release: 5.25
	Fixed in Release: 5.30
896496	Description: Fixed an issue where the network adapter configuration was not deleted in the process of the driver's uninstallation.
	Keywords: Uninstallation, network adapter configuration
	Discovered in Release: 5.25
	Fixed in Release: 5.30
897970	Description: Fixed an issue where the nd providers were not removed in the process of the driver's uninstallation.
	Keywords: Uninstallation, nd providers
	Discovered in Release: 5.25
	Fixed in Release: 5.30
832138/ 839713	Description: Fixed an issue which caused the tools' folder not to be added to the path environment variable when installing the driver on a VM.
	Keywords: mlx4_bus, Load failure
	Discovered in Release: 5.25
	Fixed in Release: 5.30
836302	Description: Fixed an issue that generated and sent the following message to the Windows event log when using firmware 2.36.5000, whenever "Mellanox WinOF Bus Counters" was selected in Perfmon: <i>"SingleFunc_196_0_0: Execution of FW command failed. op 0x68, status 0x2, errno -1"</i>
	Keywords: ECN, counters
	Discovered in Release: 5.22
	Fixed in Release: 5.30
781323	Description: Fixed an issue that occasionally caused system-hang when TCP off-load parameters were updated dynamically (in SR-IOV machines only).
	Keywords: SR-IOV machines, TCP offload parameters
	Discovered in Release: 5.25
	Fixed in Release: 5.30

Table 9 - Fixed Bugs List (Sheet 6 of 17)

Internal Ref.	Issue
831100	Description: Fixed an issue that occasionally caused system-hang upon bus driver disabling, when the encapsulation parameters were updated dynamically (in SR-IOV machines only).
	Keywords: SR-IOV machines, bus driver, encapsulation parameters
	Discovered in Release: 5.10
	Fixed in Release: 5.30
856330	Description: Updated the port type configuration as described below for adapter cards that do not support an IB port type: <ul style="list-style-type: none"> • ETH_roce_1_25:ETH_roce_1_25 for ConnectX-3 • ETH_roce_2:ETH_roce_2 for ConnectX-3 Pro
	Keywords: Port type
	Discovered in Release: 5.25
	Fixed in Release: 5.30
848108/ 848622	Description: Fixed an issue where the VF RDMA was not functional when vSwitch was attached to port 2. Now RDMA over VF is supported only when the vSwitch is attached to port 1.
	Keywords: RDMA over VF, vSwitch
	Discovered in Release: 5.25
	Fixed in Release: 5.30
861801	Description: Fixed an issue which caused the driver to hang during installation process due to a deadlock.
	Keywords: Installation, deadlock
	Discovered in Release: 5.20
	Fixed in Release: 5.30
820924	Description: Fixed an issue which limited the name to a string of 100 character length when creating Team with the Command Line.
	Keywords: Team, vlan_config, Windows Server 2008 R2
	Discovered in Release: 5.19
	Fixed in Release: 5.25

Table 9 - Fixed Bugs List (Sheet 7 of 17)

Internal Ref.	Issue
823283	Description: Fixed an issue which caused vlan_config tool on Windows Server 2008 R2 to display an error when running the showlist command on a system without any Ethernet adapters.
	Keywords: Team, vlan_config, Windows Server 2008 R2
	Discovered in Release: 5.19
	Fixed in Release: 5.25
851705	Description: Fixed an issue which occasionally caused system crash while disabling the Physical Interface in a setup where the VM switch was attached to the Mellanox interface and the vNic was created with RDMA Network Direct enabled.
	Keywords: RDMA Network Direct, vNIC
	Discovered in Release: 5.22
	Fixed in Release: 5.25
815512	Description: Fixed a lock handling on IPoIB VMQ creation error flow which caused the machine to hang due to an attempt to require a lock that was already held.
	Keywords: IPoIB VMQ, lock handling
	Discovered in Release: 5.22
	Fixed in Release: 5.25
829994/ 843730	Description: Fixed an issues which caused the bus driver to fail to start up with an old firmware (like 2.11.500) when the firmware supported a not-a-power-of-two amount of statistic counters.
	Keywords: Startup, Counter
	Discovered in Release: 5.22
	Fixed in Release: 5.25
845965	Description: Fixed an issues which caused system crash during multicast re-registrations on a setup with 2 ports both configured as IPoIB, when one of the IPoIB ports was disabled, and the SM was restarted on the second port.
	Keywords: IPoIB, SM
	Discovered in Release: 5.22
	Fixed in Release: 5.25

Table 9 - Fixed Bugs List (Sheet 8 of 17)

Internal Ref.	Issue
767937	Description: Removed the MUP registry keys for the mlx4_bus INF for Windows Server 2016.
	Keywords: Bus INF, MUP package, Windows Server 2016
	Discovered in Release: 5.19
	Fixed in Release: 5.25
824828	Description: Fixed an issue which caused the <code>SriovPort1NumVFs</code> and <code>SriovPort2NumVFs</code> fields of the <code>Get-MLNXpciDeviceSriovSetting</code> command to show the value from the registry and not the actual value set by the driver.
	Keywords: CIM, Powershell
	Discovered in Release: 5.19
	Fixed in Release: 5.25
826650	Description: Fixed an issue which caused <code>part_man</code> tool to not function on an adapter with mixed case characters.
	Keywords: <code>part_man</code> , IPoIB
	Discovered in Release: 5.19
	Fixed in Release: 5.25
834347	Description: Fixed an issue which occasionally caused BSOD when enabling VMMQ on a vport while having traffic from the VM's vport to another host. The vmswitch was created with SR-IOV enabled but the vport was not.
	Keywords: VMMQ
	Discovered in Release: 5.19
	Fixed in Release: 5.25
787139	Description: Fixed an issue that caused the <code>perf_tuning.exe</code> tool to crash in Windows Server 10 when running a dual-port configuration.
	Keywords: Tools, performance tuning, Windows Server 10
	Discovered in Release: 5.10
	Fixed in Release: 5.22
667374	Description: Old <code>SriovEnable</code> registry key is now obsolete. SR-IOV is now enabled by default.
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 9 of 17)

Internal Ref.	Issue
600519	Description: Fixed an issue which caused the system to crash due to a wrong detection of the Forwarding scenario.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22
687363	Description: Improved the notification about the port type settings events in the Windows system event Logs.
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22
685039	Description: Fixed an issue where the RDSH service had to be disabled even if the CIM provider was not selected as part of the installation.
	Keywords: Installation
	Discovered in Release: 5.10
	Fixed in Release: 5.22
691489	Description: Fixed an issue which resulted in BSOD 0x1E (KMODE_EXCEPTION_NOT_HANDLED) when the machine had multiple network port (whether a single physical card with multiple network ports or multiple single-port cards) with some ports configured with SR-IOV enabled, and some with SR-IOV disabled.
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22
742103	Description: Fixed a wrong multicast attach flow in IPoIB which caused both ports to receive multicast messages sent for one of them. Those packets were dropped later on and had no effect on the receiving traffic in the wrong destination, but it did increment the "Packets Received Error" counter.
	Keywords: IPoIB
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 10 of 17)

Internal Ref.	Issue
694482	Description: Fixed the <code>NetworkDirect API INDAadapter::Listen()</code> return status when the requested port was already in use. Changed the status from the non-compliant <code>ND_INVALID_PARAMETER</code> to the specification-compliant <code>ND_ADDRESS_ALREADY_EXISTS</code> .
	Keywords: ND/NDK
	Discovered in Release: 4.95
	Fixed in Release: 5.22
750419	Description: Added a mechanism to ensure that listeners will not receive the same ephemeral port. The driver uses the default range of ephemeral ports 49152 to 65535.
	Keywords: ND/NDK
	Discovered in Release: 5.10
	Fixed in Release: 5.22
755012	Description: Fixed an ND applications return status when a Send request failed with RNR. Changed the status from the non-compliant <code>STATUS_DATA_NOT_ACCEPTED</code> to the specification-compliant <code>STATUS_IO_TIMEOUT</code> .
	Keywords: ND
	Discovered in Release: 5.10
	Fixed in Release: 5.22
768149	Description: Fixed a race in IPoIB multicast management which caused an interface not to re-register on a multicast group after an SM restart. Such behavior resulted in traffic loss.
	Keywords: IPoIB
	Discovered in Release: 5.10
	Fixed in Release: 5.22
777590	Description: Fixed an issue which prevented the <code>'ibdiagnet --get_cable_info'</code> command from showing information on the cables in remote nodes.
	Keywords: Tools
	Discovered in Release: 4.95
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 11 of 17)

Internal Ref.	Issue
784798	Description: Fixed wrong counting of MTT-pages shown under "Mellanox WinOF Bus Counters".
	Keywords: Counters
	Discovered in Release: 5.10
	Fixed in Release: 5.22
785118	Description: Fixed an issue which caused the CIM provider installation to fail when another provider which uses the same classes was installed.
	Keywords: CIM
	Discovered in Release: 5.10
	Fixed in Release: 5.22
785384	Description: Fixed an IPoIB initialization error flow cleanup which caused a crash in the timer execution when the cleanup started before IPoIB initialization was completed successfully.
	Keywords: IPoIB
	Discovered in Release: 5.10
	Fixed in Release: 5.22
798908	Description: Fixed an issue which resulted in BSOD when the NIC received illegal <code>OID_RECEIVE_FILTER_QUEUE_PARAMETERS</code> value with QueueID higher than reported.
	Keywords: VMQ
	Discovered in Release: 5.10
	Fixed in Release: 5.22
774101	Description: Fixed an issue which prevented the TEAM interface and the VLAN interface to send and receive traffic at the same time. Now both the TEAM and the VLAN interfaces work at the same time where the TEAM interface handles all untagged frames and the VLAN interface handles the tagged frames.
	Keywords: Teaming
	Discovered in Release: 4.95
	Fixed in Release: 5.22
695898	Description: Fixed an issue causing BSOD upon driver restart in VM.
	Keywords: General
	Discovered in Release: 4.95
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 12 of 17)

Internal Ref.	Issue
681220	Description: Added to VSTAT the ability to show the ports' state in all port configurations (e.g. Ethernet, RoCE) and for several cards.
	Keywords: Tools
	Discovered in Release: 5.10
	Fixed in Release: 5.22
679691	Description: Fixed an issue that caused the system to hang when the driver was installed and RDMA applications were opened.
	Keywords: Setup
	Discovered in Release: 5.10
	Fixed in Release: 5.22
669086	Description: Disabled encapsulation offload capabilities for VFs regardless of the device capability to support it.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22
665352	Description: Added event-viewer message to identify "PortType" registry in case the port is configured for a unsupported type. For example if the device supports only Ethernet and the port is configured for IB the event-viewer message will be displayed.
	Keywords: SR-IOV
	Discovered in Release: 4.95
	Fixed in Release: 5.22
654942	Description: Fixed a bluescreen issue caused by calling NDK_FN_DISCONNECT on a not connected NDK_CONNECTOR.
	Keywords: NDK
	Discovered in Release: 5.10
	Fixed in Release: 5.22
648323	Description: Fixed an issue causing the IP fragmentation not to function properly, consequently discarding large packets, when SR-IOV was turned ON and the "Jumbo Packet" field on the "Advanced" tab of the Ethernet driver's properties (in device manager) was set to 9602 or more.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 13 of 17)

Internal Ref.	Issue
644775	Description: Modified the event viewer-log message displayed when the PCI virtualization does not support SR-IOV. The new message is: "SR-IOV cannot be enabled due to an error in the PCI_VIRTUALIZATION_INTERFACE. Possible reason, the machine does not support SR-IOV."
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22
644495	Description: Fixed an issue causing the system to crash when creating a team from the GUI.
	Keywords: Teaming, GUI
	Discovered in Release: 5.10
	Fixed in Release: 5.22
642433	Description: Fixed an issue where the TTL setting for the RDMA traffic could not be changed. The current solution is now defined as follows: 1. The driver will look for private parameter in ibbus\parameters\DefaultTTL b. If it does not exist, it will take it from TCP/IP configuration (tcpip\parameters\DefaultTTL) c. If not found, it will use TCP/IP default of 128
	Keywords: RDMA
	Discovered in Release: 5.10
	Fixed in Release: 5.22
630887	Description: Fixed an issue which resulted in interface disabling failure, and system crash when disabling the system, when multi-streams of sent fragmented/small packets could not be sent/returned to the OS.
	Keywords:
	Discovered in Release: 5.10
	Fixed in Release: 5.22
630050	Description: Fixed an issue causing a Black Screen in Windows guest OS running over non-Windows SR-IOV Hypervisor with two IPoIB adapters when the OpenSM was not running.
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 14 of 17)

Internal Ref.	Issue
623391	Description: Fixed wrong packets handling when the number of fragments equals the maximum allowed which resulted in a crash when the fragments array was higher then the allowed value.
	Keywords: Ethernet, IPoIB
	Discovered in Release: 5.10
	Fixed in Release: 5.22
623175	Description: Fixed an OpenSM issue which caused the driver to crash.
	Keywords: InfiniBand
	Discovered in Release: 5.10
	Fixed in Release: 5.22
611718	Description: Fixed an issue which caused the drivers to not work properly or crash as a result of setting values lower than the their default value.
	Keywords: General
	Discovered in Release:
	Fixed in Release: 5.22
550545	Description: Fixed an issue causing memory leakage when uninstalling the PF driver when vSwitch is connected to the device.
	Keywords: SR-IOV
	Discovered in Release:
	Fixed in Release: 5.22
587553	Description: Fixed a performance tuning issue when using the "Performance" tab in the Advanced property page. Performance tuning did not work as expected on certain systems that have processor groups enabled. Users saw unexpected RSS CPU's assigned. Now the following event log can be seen in logs: "Detected issues with Processor groups on the system that may cause performance degradation. Disable processor groups for best results."
	Keywords: Performance
	Discovered in Release: 4.95
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 15 of 17)

Internal Ref.	Issue
443600	Description: Fixed an issue where NDv2 adapter query reported MaxInlineData-Size as zero regardless of the device's capability.
	Keywords: RDMA
	Discovered in Release: 4.90
	Fixed in Release: 5.22
653552	Description: Fixed synchronization issue between NDK OIDs and driver's halt command. The driver started unloading before all the NDK OIDs were completed causing unexpected behavior in the system.
	Keywords: NDK
	Discovered in Release: 4.95
	Fixed in Release: 5.22
639974	Description: Fixed an issue that caused the system to hang when the driver failed to allocate contiguous memory on driver load.
	Keywords: General
	Discovered in Release: 5.10
	Fixed in Release: 5.22
550471	Description: Fixed wrong MAC creation from GUIDs that start with E4 D1 D2. In addition to applying the GUIDMask, the bug used to alternate the first byte of the created MAC. This alternation may have caused problems in communication with DHCP for IP address acquiring.
	Keywords: General
	Discovered in Release: 5.10
	Fixed in Release: 5.22
618647	Description: Fixed an issue where the ND applications provided with the driver package would fail to run when running on a server with WinOF-2 driver package installed.
	Keywords: RDMA
	Discovered in Release: 5.10
	Fixed in Release: 5.22
649996	Description: Fixed a rare issue where system would crash when miniport driver is reset.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 16 of 17)

Internal Ref.	Issue
669610	Description: Fixed an issue where driver could falsely indicate that the device is stuck in extreme CPU load conditions.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22
651029	Description: Fixed the issue of VF miniport reset that used to be caused by OID-timeout of OID_GEN_CURRENT_PACKET_FILTER.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22
665229	Description: Fixed RSS functionality in VM guests with SR-IOV enabled and VMs running in VMQ mode with SR-IOV enabled on host. This used to result in poor performance.
	Keywords: Ethernet
	Discovered in Release: 5.10
	Fixed in Release: 5.22
668295	Description: Fixed the issue where encapsulation settings (NVGRE+VLXAN) were not restored after reset.
	Keywords: Encapsulation
	Discovered in Release: 5.10
	Fixed in Release: 5.22
607939	Description: Fixed possible system crash when upgrading from v5.00 to a newer version.
	Keywords: WPP
	Discovered in Release: 5.10
	Fixed in Release: 5.22
623187	Description: Fixed a race between VMQ and pkey table initialization. This issue may have occurred on a setup on which a V-Switch is defined on a virtual interface (created via the part_man tool) with a non default pkey.
	Keywords: InfiniBand
	Discovered in Release: 5.10
	Fixed in Release: 5.22

Table 9 - Fixed Bugs List (Sheet 17 of 17)

Internal Ref.	Issue
625294	Description: Fixed the issue where it was possible to install the Microsoft Azure 2012R2 package on older Operating Systems.
	Keywords: Installation
	Discovered in Release: 5.10
	Fixed in Release: 5.22
666988	Description: Fixed a mismatch in the way the performance tuning and IPoIB create a MAC from a device GUID. This mismatch may have caused the performance tuning to fail finding the correct MAC for GUIDs that do not start with 0002c9, which used to result in a failure to choose the best NUMA for them.
	Keywords: IPoIB
	Discovered in Release: 5.10
	Fixed in Release: 5.22
657047	Description: Fixed an error which could cause a freeze of the drivers in the startup in a VM with several VFs.
	Keywords: SR-IOV
	Discovered in Release: 5.10
	Fixed in Release: 5.22
678607	Description: Fixed the issue where deleting a team that has VLANs attached to it would fail on Windows 10 Operating System.
	Keywords: Teaming
	Discovered in Release: 5.10
	Fixed in Release: 5.22

5 Change Log History

Table 10 - Change Log History (Sheet 1 of 28)

Category	Description	Reference Number
Rev. 5.25 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 5.25.12665. • The CIM provider version is 5.25.12665 		
Virtualization	Virtual Machine Multiple Queue (VMMQ): Virtual Machine Multiple Queues (VMMQ), formerly known as Hardware vRSS, is a NIC offload technology that provides scalability for processing network traffic of a VPort in the host (root partition) of a virtualized node. In essence, VMMQ extends the native RSS feature to the VPorts that are associated with the physical function (PF) of a NIC including the default VPort.	-
	Network Direct Kernel Provider Interface (NDK v2): The Network Direct Kernel Provider Interface (NDKPI) is an extension to NDIS that allows IHVs to provide kernel-mode Remote Direct Memory Access (RDMA) support in a network adapter.	-
	SR-IOV Port Mode: Changed the default value of Srio-vPortMode to Manual. Now, by default on dual-port devices the maximum number of VFs will be split between the two ports.	
	RDMA over VM in SR IOV Mode (Beta Level): Allows the user to work with ND and NDK over Virtual Machines when in SR-IOV mode.	
Ethernet	PacketDirect Provider Interface: PacketDirect extends NDIS with an accelerated I/O model, which can increase the number of packets processed per second by an order of magnitude and significantly decrease jitter when compared to the traditional NDIS I/O path.	-
Installation	Added support for Nano server.	
Bug Fixes	See Section 4, “Bug Fixes History”, on page 37	-
Rev. 5.22 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 5.22.12433. • The CIM provider version is 5.22.12433 		

Table 10 - Change Log History (Sheet 2 of 28)

Category	Description	Reference Number
RDMA	Improves cache hit rate in RDMA by reducing the size of the Adapter's Memory Translation Table (MTT).	-
	Changed the ND port allocation scheme from hashing to 64k bitmask.	-
	Changed the default value of RoCE mode to RoCE v2	753974
Tools	Modified the Vsat tool to function also when RoCE is disabled.	-
	nd_write_bw and nd_send_bw now support getting send completions using events instead of polling by using -e switch. Parameter only affects client side and only on "duration" mode (-D <time>).	665164
	Improved mlxtool error handling for the pkeys option. When a broken IPoIB interface registry entry exists (for example, an old virtual interface that was not fully removed), the tool would fail and exit. Now the tool skips such entries and prints the next ones.	642352
General	Enabled dual-port card to work as a single-port card.	-
Diagnostic	Improved Event Log Messages explaining behavior of the driver in case of illegal port configuration Port1: Ethernet w/o RoCE, Port2 IB.	681229
	Improved Event Log Messages issued on driver-generated dumps.	648731
SR-IOV	SR-IOV is disabled when the port is set as IB type.	-
	Set the *PriorityVLANTag registry key of Virtual Function (VF) to not support neither priority nor VLAN, and removed VlanId registry key completely.	659090
	[Beta] Added support for SR-IOV Ethernet Windows over KVM.	-
Installation	Removed the co-installer that installs performance counters. The counters is installed during the setup.	-
Teaming	Added support for tagged and untagged traffic over Team driver.	-
	Enabled teaming configuration via the Command Line Interface.	-
Documentation	Release Notes and User Manual documents were removed from the package. A new README file which includes basic installation instructions, summary of main features and requirements has replaced them.	661394

Table 10 - Change Log History (Sheet 3 of 28)

Category	Description	Reference Number
Rev. 5.19.11822 (Beta Level) Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, and IPoIB drivers version is 5.19.11822 • The CIM provider version is 5.19.11822 		
SR-IOV	Added VLAN support for NDK on VF.	
Ethernet	Added a thread race protection for RX/TX CQ/ring iterator	
Rev. 5.19.11803 (Beta Level) Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, and IPoIB drivers version is 5.19.11803 • The CIM provider version is 5.19.11803 		
NDKPI	Added support for NDKPI v2.0 interface.	
Virtualization	Added support for RoCE in SR-IOV VM. Virtualization: Added support for RoCE in virtualization mode in the hypervisor.	
PacketDirect	Added support for PacketDirect Provider Interface (PDPI).	
SR-IOV Security	Added support for the Ethertype spoof protection feature, which enables the hypervisor to control the allowed Ether-types that the VF can transmit.	
VMMQ	Added support for RSS load-balancing offload in HW for non-SRIOV VMs.	
Rev. 5.10 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 5.10.11345. • The CIM provider version is 5.10.11345 		
Operating System	Added support for a new Windows Client version - Windows 10 Client	-

Table 10 - Change Log History (Sheet 4 of 28)

Category	Description	Reference Number
General	Fixed an issue where a device state dump used for debug could cause the device to get stuck, requiring driver restart or server reboot to recover.	568240
	Fixed an issue where a system with two or more adapters could crash when one of the adapters gets disabled.	532481
	Fixed the case where during very high CPU load on the core that is in charge of transmit notification processing, driver may mistakenly decide the device is hung and initiate a reset to the interface.	550016
	Fixed the issue of when a driver fails to start (Mellanox device appears with Yellow bang in the device manager) and user stops the driver, this could cause the system to hang.	492885
General	Added third party branding for ATTO. ATTO devices will be shown in the device manager with customized ATTO device names and model numbers.	520073
RDMA	Fixed synchronization issue between client and server side in <code>nd_write_lat</code> that could cause the test to hang on start.	559668
	Fixed a problem of when a user is trying to change the RoCE mode using the "Set-MlnxDriverCoreSetting" PowerShell command on a setup with two devices, the mode does not change until the next driver restart.	520406
	Enabled the driver to use a minimal number of memory registration resources when physical memory contiguity allows it.	557943
	Implemented a resource pool to save physically contiguous memory in the driver, which was used by RDMA applications. This way the re-use of this memory will be allowed.	557954

Table 10 - Change Log History (Sheet 5 of 28)

Category	Description	Reference Number
RDMA	Improved ND connection establishment time when using firmware v2.34.5000	495620
	Changed NDK and its clients (like SMBDirect) to be disabled by default when PFC is not enabled.	449771
	Added an RDMA test to the package to demonstrate Linux interoperability.	541340
	Enlarged private data limit in ND and NDK.	562879
	Changed default RoCEv2 UDP destination port to match IANA standard.	574918
	Removed all deprecated performance tools.	569889
QoS	Fixed an issue when after OS initiated reset of the interface, default QoS policies were not properly restored.	558513
Virtualization	Fixed the issue of when enabling VMQ after SRIOV has been disabled, VMQ would not work, and the VM would fall back to paravirtualization, impacting performance.	549092
	Improved isolation of SRIOV host from VM driver issues.	549073
	Fixed a race which caused some of the configurations to be badly reinitialized during VM or host transition to VMQ mode. This bug used to cause loss of communication to the affected VM or host.	560789
	Added support in mlxtool to allow the query of PKeys configured in SR-IOV VMs.	565011
RoCE	Fixed an issue where adding VLANs would fail because the driver's internal table was not cleaned up correctly upon VLAN removal.	547762
	Fixed a memory leak caused by a race between successful finish of a Management Datagram (MAD) and canceling it.	541447
Debugging	Enabled mlxtool to allow a manual generation of register dumps.	542828
	For performance tuning purposes, debug counters were added to indicate once the driver transmit queue got full.	-
	Added a discard counter for performance analysis.	565011
	Added discard counters per Transport Class. The counters names are in beta and are subject to change.	591908

Table 10 - Change Log History (Sheet 6 of 28)

Category	Description	Reference Number
Ethernet	Fixed the adapter name in command <code>vea_man</code> to not be case sensitive,. For example: <code>vea_man -a ethernet9</code> is now the same as <code>vea_man -a Ethernet9</code> .	538603
	Fixed the issue when after removing the Virtual Ethernet Adapter, some registry keys would remain.	529621
	Enabled a device watchdog mechanism that prevents the device from sending excessive pauses to the network for any reason.	
	Added support for Windows 2008R2/Windows 7 Client teaming driver to allow selecting the MAC address of the primary interface.	514256
	Added support for driver Teaming in Windows Client 8.1.	507319
InfiniBand	Fixed the issue of when Query Path Record (QPR) option is set, a race condition occurs. The race would be between the handling of received packets and the response to the path query to the new destination. This could cause losing packets received from a new source because the path query for it was not yet finished.	536405
	Fixed the Query Path Record list handling to prevent double entering of the same destination. This bug may have caused list corruption which led to unexpected results.	535446
	Fixed the issue of when BSOD may occur when running with two HCAs and using <code>sminfo</code> when no Subnet Manager is available.	492579
	Fixed the issue of when creating virtual IPoIB PKey interfaces with HP cards using <code>part_man.exe</code> utility was not possible.	491585
	Fixed the issue where a new VM creation or its migration in IPoIB could cause the system to crash.	441213
Performance	Fixed an issue where in VMQ mode, not all receive buffers allocated for the VMQ are used, impacting performance.	567513
Installation	Added support for installing counters with co-installer. This allows the installation of counters while installing the driver via the INF mechanism.	549805
Rev. 4.95.50000 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.95.10777 • The CIM provider version is 4.95.10777 		

Table 10 - Change Log History (Sheet 7 of 28)

Category	Description	Reference Number
InfiniBand	Fixed BSOD on next driver restart when running the sminfo tool and SM is not running	492579
	Fixed instability in IPoIB driver when polling mode is enabled	521205
	Fixed the issue when live-migrated VM with virtual NIC over IPoIB physical interface loses its IP address and comes up with APIPA address (169.254.x.x)	439359
Ethernet	Fixed BSoD after the receive buffer's size changes in VMQ mode	500228/ 522073
	ECN is now configurable via PowerShell	495158
	Fixed Powershell setting of RoCE mode when the machine has more than one Mellanox device	520406
	Fixed reporting of NVGRE capabilities to the OS	535203
	Added a new mode that ignores FCS warnings and enables the Ethernet packets to be received by the NIC	-
	Added the option of setting the MAC address of 2008R2 teaming driver to be taken from the primary interface	-
	Added the option of taking the MAC Address used for teaming from one of the interfaces without manipulation	-
Changes in UI	Added the option of configuring the team MAC address to be the same as the primary adapter MAC address	514989
Troubleshooting	Added an autologger session that dumps WPP traces to file to enable easier analysis of issues	-
	Added logging of performance counters and CPU power saving state to system snapshot tool	-
RDMA	Fixed handling of send request with inline data larger than supported	443355
	Added the option to allow RDMA programmers to create pre-allocated pools of ND resources to reduce resource creation time	-
Low Level Driver	Fixed the case in which the system rarely gets unstable after driver restart	492885
Infrastructure	Updated customization for OEM cards	-

Table 10 - Change Log History (Sheet 8 of 28)

Category	Description	Reference Number
Rev. 4.90.50000 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.90.10714 • The CIM provider version is 4.90.10714 		
Generic	Fixed driver instability when handling many RDMA connection requests in parallel	461854
	Added to MLNX_System_Snapshot Mellanox specific counters and data from Get-Mlnx* Cmdlets	467529
Resiliency	Reset Flow improvements: <ul style="list-style-type: none"> • Resolved race condition when reset is initiated by more than one source • Reset initiated on one port does not cause reset of the other port 	400887
Ethernet	ETS is now configurable through DSCP values. For further details, please refer to WinOF User Manual, "Differentiated Services Code Point (DSCP)" section	434105
	Fixed the issue of when creating a Virtual Ethernet Adapter interface and removing it immediately a Blue Screen may appear	456279
	Fixed duplicated values of Receive Completion Method in Advanced Properties driver dialog on Windows Server 2012 R2	443273
	Performance Improvement: Reduced memory access time for Receive descriptors	-
	VM Scalability: More efficient handling of VMQ control path in HyperV	-
	Reduced the amount of kernel memory used for each Ethernet interface by the driver	-
	Virtual Machine traffic on the default queue now uses a single CPU core as required by Microsoft. This applies both to SRIOV and VMQ	441581

Table 10 - Change Log History (Sheet 9 of 28)

Category	Description	Reference Number
InfiniBand	Updated IBAL interface version. In order for the applications that use the IBAL interface to work with WinOF Rev 4.90.50000, they must be recompiled with the new SDK	-
	Added support for SM change event	435564
	Fixed propagation of error code when ib_join_mcast() fails	448028
	Fixed connectivity problems when using PKeys from the same partition with different membership types	417753
	Fixed VM reset after printing the message "mlx4_core 0000:05:00.0: unparavirt command: OTHER (0x3a) accepted from slave:3" in SR-IOV InfiniBand VM over non-windows hypervisor	422598
RoCE	In RoCE v2, added the option of determining the source port field of the UDP header by the application	-
NDK	Improved CPU utilization by changing ndkgetremotetokenfrommr() to return value in network byte order	-
Performance	Fixed the UI crash when working with a single port	427484
	Increased the accuracy of the run time duration parameter of ND Performance tests even when sending large message	-
	Fixed Maximum value for ThreadPoll parameter to be 200,000, instead of the previous state when it could not be set above 20,000 due to a bug	481291

Table 10 - Change Log History (Sheet 10 of 28)

Category	Description	Reference Number
IPoIB	Fixed stability issues	-
	Fixed displaying of IPoIB default turning option	428601
	Fixed temporary network connectivity issues while migrating VMs or modifying VMQ configuration for VMs that uses IPoIB with VMQ	417687
	Fixed the <code>part_man</code> tool to use the actual default <code>p_key</code> instead of <code>0xffff</code>	417858
	Fixed NIC reset when attaching to a multicast group fails	423435
	Fixed duplicated values of Receive Completion Method in UI on Windows Server 2012 R2	-
	Added support for multiple PKey interfaces in IPoIB	-
	Added support for teaming of IPoIB interfaces to allow failover	443273
	Added IPoIB adapters teaming support (beta level)	-
	Added sending of gratuitous ARP in IPoIB interface when the MAC address is changed	408388
	Reduced memory footprints of IPoIB interfaces	-
	Reduced the multiple number of path record queries to one when old query information exists	466336
	Improved completion memory access speed	440018
	Changed default VMQ/VPort affinity to use first RSS CPU	442549
	Multiple PKey support is now at GA level. The <code>part_man</code> tool allows the creation of up to 64 vIPoIB interfaces (32 per port)	-
IPoIB	Added a warning to the event log if the port MTU is higher than the reported MTU by the SM.	-
Installation	Fixed CIM failure after installation in maintenance mode	423206
	Fixed loading of old driver after driver upgrade that requires system reboot to complete the process	422812
	Fixed RoCE disable by default after installation of WinOF in Windows 8.1 Client	454020

Table 10 - Change Log History (Sheet 11 of 28)

Category	Description	Reference Number
ND	Fixed seg fault when executing ND application with no device installed or when a wrong device identifier is used	431113
	Fixed wrong reported value of supported number of SGE in 32 bit DLLs	425841
	Increased the number of supported SGEs in 32 bit DLLs to 2	425841
NVGRE	Fixed restoration of NVGRE configuration after NIC reset	442478
Changes in UI	Replaced the terms “LBFO” and “Bundle” with “Team-ing” and “team” respectively.	-
CIM/WMI	Added support to query/set/enable/disable ECN	
	Added support to query DroplessMode state	
	Fixed the issue when using the PowerShell command Get-MlnxFirmwareIdentity on a system with multiple NICs/ HCAs while one of the devices is disabled and the command fails	
Rev. 4.80.50000		
Contains the following versions of components:		
<ul style="list-style-type: none">• Bus, eth, IPoIB and mux drivers version is 4.80.10388• The CIM provider version is 4.80.10388		
Installation/ Upgrade	Added check for administrator privileges during installation	391704
	Added support for installation in silent mode without execution of perf_tune	397946
	Fixed installation stuck when Remote Desktop Session Host Windows Installer RDS compatibility is enabled	371541
Generic	Changed Reset Flow (+SR-IOV)-enabled only if no user space application is running and depends on the registry key: AllowResetOnError setting)	370536
	Changed the number of supported QPs in a multicast group from hard coded value to firmware capabilities dependent	401850
	Fixed driver load failure in machines with 1 TB memory and above	407556
	Fixed memory leak on the Virtual Machine in SR-IOV when resetting the Virtual Machine of associated VFs	373144

Table 10 - Change Log History (Sheet 12 of 28)

Category	Description	Reference Number
IPoIB	Added multiple P_Key support (beta level)	391240
	Added IPoIB SR-IOV over KVM and ESX Hypervisors (for both full and partialmembership)	-
	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	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	394703
	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	389974
	Improved transmit and receive in multi stream scenarios	-
	Enabled hardware checksum offload for non TCP/UDP traffic with ConnectX®-3 Pro	394977
	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	401655
	Fixed stability issues	-
RoCE	Added RoCE IP based	391238
ND	Fixed wrong return value in <code>IND2Adapter::QueryAddressList</code>	-
InfiniBand	Added non-default PKey in VM	-

Table 10 - Change Log History (Sheet 13 of 28)

Category	Description	Reference Number
Performance	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	414872
	Changed default perf_tuning scenario to be “Balanced configuration”	396981
	Various performance improvements	-
WMI/CIM	Added ability to read active RoCE configuration from hardware	400598
	Added support for RoCE IP Based	390573
Rev. 4.70.50050 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10143 • The CIM provider version is 4.70.10143 		
IPoIB	Fixed SM fail-over causing the driver to hang	-
Rev. 4.70.50040 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10141 • The CIM provider version is 4.70.10141 		
Generic	Optimized handling of “affinity change” on OID_RECEIVE_FILTER_QUEUE_PARAMETERS	-
	Added the ability to control the number of retries and time-out to check the device health before performing reset	-
Ethernet	Fixed missing pause response by sender when using DSCP/untag priority tag mode with ETS enabled	-
Rev. 4.70.50000 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.70.10126. • The CIM provider version is 4.70.10130 		
Installation/ Upgrade	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	-

Table 10 - Change Log History (Sheet 14 of 28)

Category	Description	Reference Number
Generic	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	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 <code>part_man</code> exit with return value 0 in case of error	-
Ethernet	Added support for Ethernet SR-IOV over Windows Hyper-V Hypervisor (over Windows Server 2012 R2)* * Requires firmware v2.30.8000 and above	-
	Added Virtual Ethernet Adapter support which enables using SMB Direct and HyperV (VMQ and NVGRE (over ConnectX®-3 Pro)) on the same port** ** Requires firmware v2.31.5050 and above.	-
	Added lossless TCP buffer management when no receive WQE are available	-
RoCE	Added ConnectX®-3 Pro support for RoCEv2	-
	Changed the transport name in <code>vstat</code> and <code>ibstat</code> to be RoCE v2.0	-
	Fixed <code>ibstat</code> 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	Fixed hard-coded limitation of 4 SGEs	-
InfiniBand	Fixed <code>vstat</code> printing of counters for Ethernet ports	-
	Fixed crash when calling <code>ib_join_mcast()</code> with <code>time-out_ms = 0</code>	330740

Table 10 - Change Log History (Sheet 15 of 28)

Category	Description	Reference Number
Performance	Improved <code>perf_tuning</code> setting in single CPU machines to avoid TX collision	-
Rev. 4.61 (Released as an intermediate release)		
Installation/ Upgrade	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	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	Added support OF IPv6 to all <code>nd_*_*</code> tests	-
	Enabled optimal interrupt moderation values in SR-IOV VF mode	-
	Stopped using <code>NdisQueryNetBufferPhysicalCount</code> to improve CPU utilization	-
IPoIB	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	-

Table 10 - Change Log History (Sheet 16 of 28)

Category	Description	Reference Number
Ethernet	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	-
Rev. 4.60.17718 Contains the following versions of components: <ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 4.60.17718. • The CIM provider version is 4.60.17718. 		
Hyper-V	Fixed NIC reset when moving IPoIB interface in a VM from non-VMQ to VMQ or from VMQ to non-VMQ	325607
Installation/ Upgrade	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	Added support for a new report for bad cables *** *** Requires firmware v2.30.8000 and above.	-
	Fixed random parsing failures of string registry entries	-
	Fixed compilation failure of “Hello_world” in the SDK	-
	Fixed the return value of ib_query_ca() if failed to allocate resources for operation	-

Table 10 - Change Log History (Sheet 17 of 28)

Category	Description	Reference Number
Performance	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	Fixed RoCE mode parsing	-
ND	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	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	-

Table 10 - Change Log History (Sheet 18 of 28)

Category	Description	Reference Number
Ethernet	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 10 - Change Log History (Sheet 19 of 28)

Category	Description	Reference Number
WMI/CIM	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 	-
Rev. 4.60.17738 Contains the following versions of components: <ul style="list-style-type: none"> • 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. 		
IPoIB	Fixed using CQ after VMQ is closed	-
	Fixed bad completion of VMQ QP that was caused by malformed WR	-
Ethernet	LBFO: Fixed the team's MAC address uniqueness in the subnet of the team in Windows Server 2008 R2	-
Rev. 4.55		

Table 10 - Change Log History (Sheet 20 of 28)

Category	Description	Reference Number
Generic	<ul style="list-style-type: none"> Added support for Windows Server 2012 R2 Operating System Added the <code>ParentBusPath</code> 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 <code>RecvCompletionMethod</code> as Interrupt is no longer supported Removed the <code>LsoV1IPv4</code> from the registry/UI Removed from the bus driver configuration the 'Non-DMA' option Removed the <code>TXRingNum</code> 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 <code>*EncapsulatedPacketTaskOff-load</code> option when using ConnectX®-3 Pro NICs 	-
Performance	<ul style="list-style-type: none"> Added the <code>nd_send_bw</code> and <code>nd_send_lat</code> ND benchmarking tools Fixed <code>nd_*_bw</code> to achieve better performance (memory buffer alignment) and consistent results 	-
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 MLX-4ND2 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 <code>part_man</code> issue related to wrong statistics over virtual partman interfaces 	-
RoCE	<ul style="list-style-type: none"> Enabled <code>roce_mode</code> value overwrite in case it exists during installation Fixed in <code>ibv_devinfo</code> 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 	-
Rev. 4.40		

Table 10 - Change Log History (Sheet 21 of 28)

Category	Description	Reference Number
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 <code>perf_tuning</code> failure during the installation. An error message will be printed in the installation log upon <code>perf_tuning</code> 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 <code>TXRingNum</code> Changed the <code>perf_tuning</code> setting to achieve a better performance tuning Added the <code>nd_write_bw/nd_write_lat</code> and <code>nd_read_bw/nd_read_lat</code> tools Fixed the <code>perf_tuning</code> indication of the last chosen tuning scenarios Fixed a crash in the <code>ib_send_lat/bw</code> utilities caused when the port link was down Fixed the “Restore to defaults” option in the <code>perf_tuning</code> tool. Now the default values are being restored 	-

Table 10 - Change Log History (Sheet 22 of 28)

Category	Description	Reference Number
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 In addition to old values that are not supported anymore. 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 	-

Table 10 - Change Log History (Sheet 23 of 28)

Category	Description	Reference Number
Ethernet	<ul style="list-style-type: none"> 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 End-Point 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 End-Point 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 In addition to old values that are not supported anymore. 	-

Table 10 - Change Log History (Sheet 24 of 28)

Category	Description	Reference Number
Ethernet	<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 End-Point 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: 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 	

Table 10 - Change Log History (Sheet 25 of 28)

Category	Description	Reference Number
InfiniBand	<ul style="list-style-type: none"> On rare occasions, depends on the GUID assignment, the IPoIB MAC address can be assigned with a multi-cast 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 <code>active_mtu</code> field to <code>struct ib_port_attr_t</code> Added the option of <code>vstat</code> displaying the <code>active_mtu</code> 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 <code>ibv_devinfo</code> displaying the correct MTU value after it was changed Added the option of <code>part_man</code> printing the adapter name when the Port GUID is set to zero. Added the option of <code>part_man</code> 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 <code>portsetting</code> registry key during uninstallation 	
Rev. 4.3 (This version was released as an intermediate release)		
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 Server 2012 are: <ul style="list-style-type: none"> <code>RssMaxProcNumber</code> <code>NumRSSQueues</code> <code>RSSProfile</code> 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 	

Table 10 - Change Log History (Sheet 26 of 28)

Category	Description	Reference Number
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 	
Rev. 4.2		
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 	

Table 10 - Change Log History (Sheet 27 of 28)

Category	Description	Reference Number
Generic	<ul style="list-style-type: none"> 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 10 - Change Log History (Sheet 28 of 28)

Category	Description	Reference Number
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	