



Mellanox WinOF VPI Release Notes

Rev 5.50

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “ASIS” 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 2018. 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®, NPS®, 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 Supported Operating System Versions	7
1.3 Hardware and Software Requirements	8
1.4 Certifications	8
1.5 Supported Network Adapter Cards	9
1.5.1 Firmware Versions	9
1.5.2 MFT Version	9
1.6 RoCE Modes Matrix	9
Chapter 2 Changes and New Features in Rev 5.50	11
2.1 Beta Features	12
2.2 Unsupported Functionalities/Features	12
Chapter 3 Known Issues	13
Chapter 4 Bug Fixes History	16
Chapter 5 Change Log History	29
Chapter 6 API Change Log History	51

List of Tables

Table 1:	Release Update History	5
Table 2:	Supported Operating System Versions	7
Table 3:	Certifications	8
Table 4:	Supported Network Adapter Cards	9
Table 5:	Firmware Versions	9
Table 6:	MFT Version	9
Table 7:	RoCE Modes Matrix	9
Table 8:	RoCE v2 UDP Port Matrix	10
Table 9:	Changes and New Features in Rev 5.50	11
Table 10:	Beta Features	12
Table 11:	Unsupported Functionalities/Features	12
Table 12:	Known Issues	13
Table 13:	Fixed Bugs List	16
Table 14:	Change Log History	29
Table 15:	API Change Log History	51

Release Update History

Table 1 - Release Update History

Release	Date	Description
Rev 5.50	October 03, 2018	Added support for Windows Server 2019, thus updated the following sections: <ul style="list-style-type: none"> • 1.2 “Supported Operating System Versions,” on page 7: Added Windows Server 2019 where applicable • 1.4 “Certifications,” on page 8: Added Windows Server 2019 • 2 “Changes and New Features in Rev 5.50,” on page 11: Added explanation on the RSC behavior when using Windows Server 2019
	August 13, 2018	Updated Table 11, “Unsupported Functionalities/Features,” on page 12
	May 15, 2018	Windows 10 Client 1803 (RS4) is now GA and certified.
	May 06, 2018	Initial release of WinOF Rev 5.50. This version introduces new changes and features (see Section 2, “Changes and New Features in Rev 5.50,” on page 11) and bug fixes (Section 4, “Bug Fixes History,” on page 16).

1 Overview

These are the release notes for the release of Mellanox WinOF VPI Drivers Rev 5.50 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 5.50 for Windows package contains the following components:

- Core and ULPs:
 - IB HCA low-level drivers (mlx4)
 - Ethernet driver (ETH)
 - IP over InfiniBand (IPoIB)
 - Network Direct (ND)
 - Network Direct Kernel (NDK) Provider Interface
 - MUX driver for IPoIB and Client OSes
- 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
 - mlxtool
 - 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 Supported Operating System Versions

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

Table 2 - Supported Operating System Versions

Supported Host OS	Supported Guest OS
Virtualization Mode: None	
Windows Server 2012	N/A
Windows Server 2012 R2	N/A
Windows Server 2016	N/A
Windows Server 2019	N/A
[Beta] Windows Server 1803 (RS4)	N/A
Windows 8.1 Client (64 bit only)	N/A
Windows 10 Client (64 bit only)	N/A
Windows 10 Client 1803 (RS4) (64 bit only)	N/A
Virtualization Mode: VMQ	
Windows Server 2012	Any supported guest OS for Hyper-V
Windows Server 2012 R2	Any supported guest OS for Hyper-V
Windows Server 2016	Any supported guest OS for Hyper-V
Windows Server 2019	Any supported guest OS for Hyper-V
[Beta] Windows Server 1803 (RS4)	Any supported guest OS for Hyper-V
Virtualization Mode: VMMQ	
Windows Server 2016	Any supported guest OS for Hyper-V
Windows Server 1803 (RS4)	Any supported guest OS for Hyper-V
Windows Server 2019	Any supported guest OS for Hyper-V
Virtualization Mode: Hyper-V (SR-IOV)	
Windows Server 2012 R2	Windows Server 2012 Windows Server 2012 R2 Windows Client 8.1 Windows Client 10
Windows Server 2016	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 Windows Client 8.1 Windows Client 10
[Beta] Windows Server 1803 (RS4)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 Windows Server 1803 (RS4) Windows Client 8.1 Windows Client 10

Table 2 - Supported Operating System Versions

Supported Host OS	Supported Guest OS
Windows Server 2019	Windows Server 2019 Windows Server 2016
SR-IOV InfiniBand KVM	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016
SR-IOV Ethernet KVM	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016

1.3 Hardware and Software Requirements

The following are the hardware and software requirements of WinOF Rev 5.50:

- The maximum number of supported CPUs is:
 - up to 252 logical processors when Hyper-threading is enabled
 - up to 126 logical processors when Hyper-threading is disabled
- Administrator privileges on your machine(s)
- Disk space: Minimum 100MB

1.4 Certifications

The following describes the driver’s certification status per operating system.

Table 3 - Certifications

Operating System	Logo Certification	SDDC Premium Certification
Windows 8.1 Client	Certified	N/A
Windows 10 Client	Certified	N/A
Windows 10 Client 1803 (RS4)	N/A	Certified
Windows Server 2012	Certified	N/A
Windows Server 2012 R2	Certified	N/A
Windows Server 2016	Certified	N/A
Windows Server 1803 (RS4)	N/A	N/A
Windows Server 2019	N/A	N/A

1.5 Supported Network Adapter Cards

Mellanox WinOF Rev 5.50 supports the following Mellanox network adapter cards:

Table 4 - Supported Network Adapter Cards

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

a. This version is supported by the ConnectX®-3 adapter card, but is officially qualified for ConnectX®-3 Pro only.



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

1.5.1 Firmware Versions

Mellanox WinOF Rev 5.50 provides the following firmware for Mellanox NICs:

Table 5 - Firmware Versions

NICs	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3 Pro / ConnectX®-3 Pro EN	2.42.5000	2.40.7000
ConnectX®-3 / ConnectX®-3 EN	2.42.5000	2.40.7000

1.5.2 MFT Version

Mellanox WinOF Rev 5.50 provides the following MFT version:

Table 6 - MFT Version

	Supported Version
MFT for Windows	4.9.0

1.6 RoCE Modes Matrix

The following is RoCE modes matrix.

Table 7 - 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)

Table 7 - 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
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 8 - 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 5.50



This package version is 5.50.50000 and 5.50.50010 for Windows Server 2019. The package contains the following versions of components:

- Bus, eth, IPoIB and MUX drivers version is 5.50.14643
- The CIM Provider version is 5.50.14643

Table 9 - Changes and New Features in Rev 5.50

Category	Description
Dump Me Now (DMN)	DMN is a bus driver (mlx4_bus.sys) feature that generates dumps and traces from various components, including hardware, firmware and software, upon internally detected issues (by the resiliency sensors), user requests (mlxtool) or ND application requests via the extended Mellanox ND API. DMN is unsupported on VFs. For further information, refer to section <i>Dump Me Now (DMN)</i> in the User Manual
CPUs	Added supports for systems with up to 252 logical processors when Hyperthreading is enabled and up to 126 logical processors when Hyperthreading is disabled.
Performance	Added support for RSC solution in TCP/IP traffic to reduce CPU overhead.
	Added support for extended NDSPI to control CQ moderation.
	Added a new counter for packets with no destination resource.
	Added a new registry key that allows users to configure the E2E Congestion Control feature.
	Added to the vlan_config tool the ability to create VLANs for the Physical Function (PF) in addition to the Virtual Function (VF).
IPoIB	Added support for VMQ over IPoIB in Windows Server 2016.
Debugging	Added support for collecting firmware MST dumps in cases of system bug check.
	Added an event log message (ID 273) that is printed when the number of resources to load the VF is insufficient.
Counters	Added a counter for the number of packets discarded due to an invalid QP number.
	Added DSCP based counters to support traffic where no VLAN/priority is present.
RDMA	RoCE TTL default value has been changed from 128 to 16.
General	Added support for Windows Server 2019. Note: Currently the drivers in the Windows Server 2019 package are certified only for Windows Server 2016.
RSC	Modified the RSC default mode when using Windows Server 2019. RSC is disabled by default in Windows Server 2019.
Bug Fixes	See “Bug Fixes History” on page 16 .

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

2.1 Beta Features

Table 10 - Beta Features

Category	Description
ibdump	ibdump is currently at beta level.
IPoIB	IPv6 support of IPoIB in an SR-IOV guest OS over KVM
	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 13.

2.2 Unsupported Functionalities/Features

The following are the unsupported functionalities/features in WinOF:

Table 11 - Unsupported Functionalities/Features

Functionality/Feature	Not Supported as of WinOF Version
Software RSC for tunneled traffic	WinOF v5.50
Wake-On-Lan	
RDMA in the Guest OSes	
ND over virtual switch attached to IPoIB port	
Memory Translation Table (MTT) Optimization	
RDMA over VM when in SR IOV mode	WinOF 5.35
IBVerbs	WinOF 5.22
WinVerbs	
IBAL performance tools (ib* ibv*)	WinOF 5.10
IBAL ND provider	
IPoIB team configuration through PowerShell	WinOF 4.90
IPv6 over IPoIB team ports	
VMQ over IPoIB team ports	
ConnectX®-2 adapter cards	
ND over WinVerbs provider	WinOF 4.52
SRP	

3 Known Issues

Table 12 provides a list of known bugs and limitations in regards to this release of WinOF.

For a list of old Known Issues, please see the [WinOF_Archived_Known_Issues](#) file.

Table 12 - Known Issues (Sheet 1 of 3)

Internal Ref.	Issue
1371587	Description: When in IPoIB mode, changing the receive buffers size using a registry key and restarting the driver while the RDMA traffic is running may result in a command failure.
	Workaround: N/A
	Keywords: IPoIB
	Detected in version: 5.50
1368272	Description: The “Disable-NetAdapterRDMA” command disables the NDK activity only, ND activity is not affected by it.
	Workaround: N/A
	Keywords: NDK, ND
	Detected in version: 5.50
1297888	Description: On Windows 2016 Client: The SR-IOV virtual adapters' counters do not count the data of the Virtual Function. The counters will not rise when the Virtual Function's counters rise neither in the performance monitor nor in the task manager.
	Workaround: N/A
	Keywords: SR-IOV, VF Counters
	Detected in version: 5.50
1139573	Description: The maximum value of the *NumRSSQueues registry key is 64.
	Workaround: N/A
	Keywords: *NumRSSQueues registry key
	Detected in version: 5.50
1048287	Description: When setting the MaxCMRetries value to 1, the REQ, REP, or DREQ messages are not sent.
	Workaround: N/A
	Keywords: MaxCMRetries
	Detected in version: 5.50
1158964	Description: The driver will ask for a reboot when a network interface gets disconnected while loading the VMs on hyper-V.
	Workaround: Disable and enable the driver manually, without rebooting. Note, you may have to perform this action twice to resolve the issue.
	Keywords: VMs, hyper-V, reboot
	Detected in version: 5.50
1158964	Description: When a network interface gets disconnected while loading VMs on hyper-V, the driver will ask for reboot.
	Workaround: Disable and enable the driver manually, without rebooting. You may have to perform this action twice to resolve the issue.
	Keywords: VMs, hyper-V, reboot
	Detected in version: 5.50

Table 12 - Known Issues (Sheet 2 of 3)

Internal Ref.	Issue
1130716	Description: The <code>mlxtool show devices</code> command is not supported for VFs.
	Workaround: N/A
	Keywords: <code>mlxtool show devices</code> ; VFs
	Detected in version: 5.50
1046418	Description: Switching between upgrading with INF and upgrading with the MSI package can end up with a different version of the Ethernet/IPoIB driver and the bus driver.
	Workaround: Use the same upgrade methodology in all upgrades. Following the upgrade, update the device driver via the device manager.
	Keywords: Installation, INF only installation.
	Detected in version: 5.40.54000
928517	Description: When configuring certain device settings to invalid values in the driver's Advanced Properties tab, a random number is used as the actual value, instead of the default number.
	Workaround: N/A
	Keywords: device settings, advanced properties
	Discovered in Release: 5.40.54000
1034664	Description: In case the network adapter of the remote ND connection was reset, the local peer does not receive a disconnect notification.
	Workaround: To make sure that the connection is still valid, post a send request on the QP and check its completion status.
	Keywords: ND, disconnect notification
	Discovered in Release: 5.40.54000
1036061	Description: When attempting to allocate several VFs to several VMs, in some cases the VF will not be allocated to the VM, and the VM will continue to work on the synthetic path. The user can identify the issue by running the <code>Get-Netadaptervport</code> command, and making sure that the number of MAC addresses is equal to the number of VFs. The issue occurs due to a bug in the Windows NetVSC co-installer, and has been fixed in a QFE. It occurs only in Windows Server 2012 R2.
	Workaround: Run the following commands in the problematic VM: <pre>RUNDLL32.exe pnpclean.dll,RunDLL_PnpClean /devices /maxclean netcfg -c p -i netvsc_vfpp -l c:\windows\inf\wnetvsc_vfpp.inf</pre>
	Keywords: SR-IOV
	Discovered in Release: 5.40.54000
1020140	Description: When using Windows Server 2012 R2 or Windows Server 2016, in some cases a system crash may occur due to an OS bug when disabling the NDK. The issue will be resolved in the next version of Windows Server.
	Workaround: N/A
	Keywords: NDK, BSOD
	Discovered in Release: 5.40.54000

Table 12 - Known Issues (Sheet 3 of 3)

Internal Ref.	Issue
981757	<p>Description: After driver upgrade, The ND state might be invalid, and the following event might appear in the event viewer: "Ndfldr: ND is in invalid state as a result of a mismatch between the ndfltr.sys driver version and mlx4_bus.sys driver version."</p> <p>Workaround: Close all applications that use ND before upgrading the driver or upgrade the driver first, close all applications that use ND, and restart the bus driver.</p> <p>Keywords: Installation, upgrade, ND</p> <p>Discovered in Release: 5.40.54000</p>
1081254	<p>Description: RDMA Activity counters are not support in IPoIB.</p> <p>Workaround: N/A</p> <p>Keywords: RDMA Activity counters, IPoIB</p> <p>Detected in version: 5.40</p>
-	<p>Description: When an Ethernet port is set to SR-IOV, InfiniBand is not supported on the second port.</p> <p>Workaround: N/A</p> <p>Keywords: SR-IOV, Ethernet, InfiniBand</p> <p>Detected in version: 5.40</p>
1337021	<p>Description: Received RDMA Activity Ack packets are counted as 20B instead of ~70B.</p> <p>Workaround: N/A</p> <p>Keywords: Counters</p> <p>Detected in version: 5.40</p>

4 Bug Fixes History

Table 13 lists the bugs fixed in this release. For a list of old Bug Fixes, please see [WinOF_Archived_Bug_Fixes](#). file.

Table 13 - Fixed Bugs List (Sheet 1 of 13)

Internal Ref.	Issue
1365213	Description: Fixed an issue that occasionally caused the system to crash when the vNic was detached from the VM during heavy traffic when in VMQ\VMMQ mode.
	Keywords: VMQ\VMMQ mode
	Discovered in Release: 5.40.54000
	Fixed in Release: 5.50
1344714	Description: Fixed an issue where the RoCE connection failed as a result of an incorrect GID when the Universal/Local (U/L) bit in the MAC was set to 1.
	Keywords: RoCE
	Discovered in Release: 5.40.54000
	Fixed in Release: 5.50
1042285/ 1042206	Description: Fixed an issue that caused the mlxtool PDDR tool to provide some inaccurate information for Infiniband links.
	Keywords: mlxtool, PDDR, IPoIB
	Discovered in Release: 5.40.54000
	Fixed in Release: 5.50
1201166	Description: Disabled the option to stop the uninstall process once the driver uninstallation process started.
	Keywords: Driver uninstallation process
	Detected in version: 5.40.54000
	Fixed in Release: 5.50
1327365	Description: Fixed an issue that caused networks with new Subnet Managers (OpenSM 4.7.0 and up) to drop malformed multicast-join packets issued by the driver. The driver now constructs the multicast join request correctly.
	Keywords: OpenSM, multicast packets
	Detected in version: 5.35
	Fixed in Release: 5.50
1340828	Description: In case the DSCP values are lower than the max priority i.e: DSCP(4)->Prio(0) when mapping the DSCP to a certain priority, the priority's value will be set the same as the DSCP's value.
	Keywords: DSCP counters
	Detected in version: 5.40
	Fixed in Release: 5.50

Table 13 - Fixed Bugs List (Sheet 2 of 13)

Internal Ref.	Issue
1333906	Description: Fixed an issue that caused the driver to hang when issuing an <code>OID_SRIO-V_RESET_VF</code> request to reset a specified PCI Express (PCIe) Virtual Function (VF) due to a race between the resiliency flow and the FLR request.
	Keywords: Driver hang, PCI Express (PCIe) Virtual Function (VF)
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1252614	Description: Fixed an issue that caused the driver to reset the adapter as a result of a false alarm of a stuck receive queue.
	Keywords: False alarm, Receive queue
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1266230	Description: Fixed an issue that caused a Black Screen upon driver's removal due to extremely low memory conditions, when the memory allocations started to fail.
	Keywords: RDMA, NDK, Black Screen
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1261837	Description: Fixed an issue that caused the binding to overrun the ND function <code>INDEndpoint</code> error status when it returned from the underlying functions. This resulted in wrong status display of the MR. The MR was displayed as registered when it was not, thus prevented the user from accessing it. This fix verifies that the user will receive the correct error status upon such scenario.
	Keywords: ND function <code>INDEndpoint</code> , MR
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1284856	Description: Fixed an issue that limited the number of MSI-X cores in Virtual Function to 8. Now the limited the number of MSI-X is 128 cores.
	Keywords: MSI-X vectors, VFs
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1134253	Description: Fixed a BSOD that occurred on servers with more than 64 cores as the Tx traffic did not honor the Tx affinity implied by the TSS, when the number of potential RSS CPUs was greater than 64.
	Keywords: Tx traffic, RSS CPUs, TSS, BSOD
	Discovered in Release: 5.35
	Fixed in Release: 5.50
1078811	Description: When the <code>mlxtool dbg resources</code> command is executed, the <code>FS_RULE</code> quota number is displayed instead of the "Managed by PF" message.
	Keywords: <code>mlxtool</code> , <code>dbg resources</code>
	Detected in version: 5.30
	Fixed in Release: 5.50

Table 13 - Fixed Bugs List (Sheet 3 of 13)

Internal Ref.	Issue
1081576	<p>Description: When setting the LogNumQp and LogNumRdmaRc registry settings to their maximum value, the WinOF bus driver fails to load.</p> <p>Keywords: LogNumQp, LogNumRdmaRc, driver load failure</p> <p>Detected in version: 5.40</p> <p>Fixed in Release: 5.50</p>
1079136	<p>Description: The "TX Ring Is Full Packets" perfmon counter is not functioning properly on IPoIB.</p> <p>Keywords: Perfmon counter, IPoIB</p> <p>Detected in version: 5.40</p> <p>Fixed in Release: 5.50</p>
1038193	<p>Description: When installing the driver over Windows 2012R2 inbox driver, the Log-NumQP parameter remains in the registry. Thus, a number of QPs are limited to 64K instead of 512K (the driver's default).</p> <p>Keywords: Windows 2012R2 inbox, LogNumQP</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1243974	<p>Description: Fixed an issue that caused a system crash during driver startup when working in RSS mode.</p> <p>Keywords: RSS, system crash, driver startup</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1170913	<p>Description: Fixed an issue that caused a system crash when the interface connected to vSwitch was disabled and the operating system did not clean all VMQs.</p> <p>Keywords: IPoIB, VMQ</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1213675	<p>Description: Fixed a race in the Communication Manager that could cause a crash while attempting to obtain ND/NDK connection. This was an atypical issue that required an unusual timing of events.</p> <p>Keywords: CM, Communication Manager, connection, ND, NDK</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1081160	<p>Description: Fixed an issue that caused commands failure and protection domain violation when running the ND application.</p> <p>Keywords: ND application, commands failure, protection domain violation</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>

Table 13 - Fixed Bugs List (Sheet 4 of 13)

Internal Ref.	Issue
1170202	Description: Fixed an issue in the mlxtool, where the “mlxtool dbg ipoib-ep [<Interface Name>]” command reported partial results of the EndPoint list when there was a large number of endpoints.
	Keywords: mlxtool, dbg ipoib-ep
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1172093	Description: Fixed an issue that caused the VM to crash when restarting the PF drivers and their peers in the target machine.
	Keywords: PF, VF, driver restart, port down
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1182951	Description: Fixed an issue that caused a memory leak when RoCE was enabled.
	Keywords: Memory leak, RoCE
	Discovered in Release: 5.30
	Fixed in Release: 5.50
1190576	Description: Fixed an issue that set a wrong value to the *ReceiveBuffers key when it was restored to default.
	Keywords: INF,*ReceiveBuffers
	Discovered in Release: 5.30
	Fixed in Release: 5.50
1065413	Description: Fixed a crash that occurred when changing the Ethernet IP address while RDMA traffic was running.
	Keywords: Crash, IP address change, RDMA traffic
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1070844	Description: Fixed a crash that occurred on IPoIB driver stack.
	Keywords: Crash, driver teardown, IPoIB
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1082383	Description: Fixed a BSOD that occurred when a memory allocation failed upon driver startup.
	Keywords: Driver load, memory allocation failure, BSOD
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1168384/ 1170019	Description: Fixed an issue where the connection port numbers did not increase sequentially when running nd_*_bw application with multiple QPs.
	Keywords: Connection port numbers, nd_*_bw
	Discovered in Release: 5.40
	Fixed in Release: 5.50

Table 13 - Fixed Bugs List (Sheet 5 of 13)

Internal Ref.	Issue
1131583	Description: Fixed an issue that caused a crash upon ND connection establishment.
	Keywords: CEP, Blue Screen crash
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1134253	Description: Fixed an issue where the Tx traffic did not honor the Tx affinity implied by the TSS when the number of potential RSS CPUs was greater than 64.
	Keywords: TSS, RSS
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1038193	Description: Fixed an issue that limited the number of QPs to 64K instead of 512K (the driver's default) when installing the driver over Windows 2012 R2 Inbox driver as the Log-NumQP parameter remained in the registry/
	Keywords: Windows 2012R2 inbox, LogNumQP
	Discovered in Release: 5.40
	Fixed in Release: 5.50
936651	Description: Fixed an issue where removing a PKey that was a part of an IPoIB team interface disabled the team and the option to delete it.
	Keywords: IPoIB Pkeys, Team
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1117581	Description: Added support for servers with more than 64 cores.
	Keywords: VMQ, SR-IOV
	Discovered in Release: 5.30
	Fixed in Release: 5.50
1037915	Description: Fixed a deadlock in which the driver could get into an error state in case resetting it and attempting to connect to it were performed simultaneously.
	Keywords: deadlock, driver reset
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1077015	Description: Fixed an issue that could cause the <code>Get-MlnxPCIDeviceSriovSetting</code> command to display a wrong number of HCAs.
	Keywords: <code>Get-MlnxPCIDeviceSriovSetting</code> , HCA number
	Discovered in Release: 5.45
	Fixed in Release: 5.50
1078808	Description: Fixed an issue where the <code>mlxtool dbg resources</code> command failed to pull information about the last VF, and showed the PF as VF0.
	Keywords: <code>mlxtool</code>
	Discovered in Release: 5.40
	Fixed in Release: 5.50

Table 13 - Fixed Bugs List (Sheet 6 of 13)

Internal Ref.	Issue
1081045	<p>Description: Fixed an issue where using invalid parameters in <code>mlxtool perfstat</code> command leads to an infinite waiting time.</p> <p>Keywords: <code>mlxtool</code>, <code>perfstat</code></p> <p>Discovered in Release: 5.45</p> <p>Fixed in Release: 5.50</p>
1117215	<p>Description: Fixed an issue where the <code>Get-MlnxPCIDeviceSriovSetting</code> command failed on a server with more than one device, when one of the devices was disabled. Following the fix, the command returns results only for the devices that are up.</p> <p>Keywords: CIM, SR-IOV</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1118060	<p>Description: Fixed an issue that caused an excessively long installation time when installing the driver on Windows Server 2012 virtual machine in SR-IOV mode with more than 2 VFs.</p> <p>Keywords: Installation, virtual machine, VM</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1131167	<p>Description: Fixed an issue that could cause a wrong link state when the PF physical port was 2.</p> <p>Keywords: SR-IOV, VF</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1050738	<p>Description: Fixed the issue of when running <code>mlxtool show perfstats</code>, incorrect Rx counters were returned when working in polling mode.</p> <p>Keywords: <code>perfstats</code>, <code>mlxtool</code></p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
1150078	<p>Description: Fixed a memory leakage that occurred upon requests of 16 million QPs.</p> <p>Keywords: Memory leakage, configuration, resources</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>
946432	<p>Description: Fixed an issue that turned the vSwitch port to internal when a bus driver restart was followed by a miniport driver restart.</p> <p>Keywords: vSwitch, IPoIB</p> <p>Discovered in Release: 5.40</p> <p>Fixed in Release: 5.50</p>

Table 13 - Fixed Bugs List (Sheet 7 of 13)

Internal Ref.	Issue
989781	Description: Fixed an issue that could cause a black screen on a driver startup in a VM with two VFs or more.
	Keywords: VM, VF, black screen, SR-IOV
	Discovered in Release: 5.40
	Fixed in Release: 5.50
1297549	Description: Fixed a BSOD that occurred while installing WinOF v5.35 due to stack usage overrun.
	Keywords: BSOD, stack usage overrun
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
909274/ 896999	Description: RoCE fallback mode is not working when SR-IOV is enabled.
	Keywords: RoCE, fallback mode, SR-IOV
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
964757	Description: On servers where Hyper-v role is installed and SR-IOV is enabled, killing the ibdump process will cause a resource leak.
	Keywords: ibdump, resource leak
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
1081160	Description: Fixed an issue that caused commands failure and protection domain violation when running the ND application.
	Keywords: ND application, commands failure, protection domain violation
	Discovered in Release: Exists from day one
	Fixed in Release: 5.40.54000
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.
	Keywords: hypervisor, vmswitch, SR-IOV
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
1064443	Description: Fixed an issue that could cause a system crash on driver load in rare cases. This could occur when the driver was waiting for firmware boot to be completed before accessing the firmware.
	Keywords: system crash, driver load, firmware boot
	Discovered in Release: 5.30
	Fixed in Release: 5.40 54000
1059536	Description: Fixed an issue that allowed executing the QP state change command when the QP was not in a valid state. This caused event viewer flooding.
	Keywords: QP state change, event viewer
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000

Table 13 - Fixed Bugs List (Sheet 8 of 13)

Internal Ref.	Issue
1064443	Description: Fixed an issue that could cause a system crash on driver load in rare cases. This could occur when the driver was waiting for firmware boot to be completed before accessing the firmware.
	Keywords: system crash, driver load, firmware boot
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
946432	Description: Fixed an issue that turned the vSwitch port to internal when a bus driver restart was followed by a miniport driver restart.
	Keywords: vSwitch, IPoIB
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
1020562	Description: Added an event log when a lost EQE interrupt is detected (event ID 156).
	Keywords: Event log, lost EQE interrupt
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
1007026	Description: Fixed an issue where new ND connections cannot be established while ibdump tool is running.
	Keywords: IBDump, ND, RDMA
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
1038291	Description: Fixed an issue that caused the ibstat tool to report the wrong link speed. The issue occurred on Ethernet ports when the link speed on the switch was set to 1G and RoCE was enabled.
	Keywords: ibstat, Link Speed, 1G
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
987803	Description: Fixed an issue that caused a failure in opening new ND/NDK connections.
	Keywords: ND, NDK, RDMA
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
882140	Description: Fixed the IPoIB network interface to show the link's signaling rate.
	Keywords: IPoIB, signaling rate
	Discovered in Release: 5.02
	Fixed in Release: 5.40.54000
866178	Description: Fixed an issue where disabling the driver could cause a hang when opensm.exe was running on the machine.
	Keywords: MAD, IBAL, IBBUS
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000

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

Internal Ref.	Issue
964590	Description: Fixed an issue where the VF can remain in an unclear state in case of reset during the loading phase.
	Keywords: VF, SR-IOV
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
987804	Description: Fixed an issue where RDMA applications could hang following a miniport restart.
	Keywords: RDMA, miniport restart
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
800647	Description: Fixed an issue where during a high CPU load the driver could mistakenly detect a device hang, and issue a NIC reset.
	Keywords: High CPU load, device hang, NIC reset
	Discovered in Release: 5.20
	Fixed in Release: 5.40.54000
584926	Description: Fixed a memory leak in the mlx4u and mlx4nd dll files.
	Keywords: ND, mlx4u, mlx4nd, memory leak
	Discovered in Release: 5.00
	Fixed in Release: 5.40.54000
676863	Description: Fixed an issue that could lead to a high CPU utilization. Following the fix, packets with unrecognized Ethernet protocol are dropped and an appropriate message is printed to the event log.
	Keywords: IPoIB, high CPU utilization
	Discovered in Release: 4.90
	Fixed in Release: 5.40.54000
1005508	Description: Fixed an issue where an ND call to the CancelOverlappedRequests() (Interface IND2Overlapped) function returned an incorrect return value. The fix correctly returns ND_SUCCESS instead of the incorrect ND_PENDING in case of a successful function call.
	Keywords: ND application
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
1005018	Description: Fixed an issue that caused ND application crash when creating a shared receive queue, and overloading the queue with post receives.
	Keywords: ND Application, SRQ
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
967654	Description: Fixed an issue that could lead to BSOD when removing a Pkey.
	Keywords: Blue screen, IPoIB, Pkey
	Discovered in Release: 5.20
	Fixed in Release: 5.40.54000

Table 13 - Fixed Bugs List (Sheet 10 of 13)

Internal Ref.	Issue
1022250	Description: Fixed an issue that could cause BSOD when resizing a number of Queue Pairs (QPs) in Virtual Multiple Machine Queue (VMMQ) mode, while running stress traffic to the VM.
	Keywords: Blue screen, VMMQ
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
1022033	Description: Fixed an issue where the internal tracing mechanism could cause driver memory corruption during high stress of commands to the firmware, while writing debug information to the file.
	Keywords: Internal tracing, driver memory corruption, debug information
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
1029795	Description: Fixed an issue that could cause a memory leak in the bus driver following multiple resets.
	Keywords: Mlx4_Bus, memory leak
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
884771	Description: Fixed an issue where wrong driver hang detection could lead to a miniports reset.
	Keywords: driver hang, mini-ports, reset
	Discovered in Release: 5.30
	Fixed in Release: 5.40.54000
974818	Description: Fixed an issue that wrongly showed RoCE v1 instead of No RoCE as a transport mode in the virtual machine VSTAT output.
	Keywords: VM, guest
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
933278/ 935344	Description: Removed the following redundant VM Event Log messages: 122, 104, 144.
	Keywords: VM, Guest
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000
1000565	Description: Fixed an issue that could cause kernel memory leakage in the Ethernet driver.
	Keywords: Ethernet
	Discovered in Release: 5.02
	Fixed in Release: 5.40.54000
940765	Description: Fixed a wrong Link State value in the VSTAT.
	Keywords: VSTAT
	Discovered in Release: 5.35
	Fixed in Release: 5.40.54000

Table 13 - Fixed Bugs List (Sheet 11 of 13)

Internal Ref.	Issue
572934	Description: Fixed an issue where the "mlxtool dbg mstdump" command wrongly returned success value, in case the MST dump feature was disabled.
	Keywords: mlxtool , mstdump
	Discovered in Release: 5.10
	Fixed in Release: 5.40.54000
980191	Description: Fixed an issue that could cause a system crash in a shutdown scenario.
	Keywords: NDIS, system crash, shutdown
	Discovered in Release: 5.40
	Fixed in Release: 5.40.54000
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: 5.35.52000 (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: 5.35.52000 (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 13 - Fixed Bugs List (Sheet 12 of 13)

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
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

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

Internal Ref.	Issue
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

5 Change Log History

Table 14 - Change Log History (Sheet 1 of 22)

Category	Description	Reference Number
Rev. 5.40		
Contains the following versions of components:		
<ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 5.40.13749 • The CIM provider version is 5.40.13749 		
General	Added PDDR (Port Diagnostics Database Register) support: providing details on the root cause when the link is down via mlxtool.	956728
	Improved driver load time by reducing the amount of context initialization work done in the firmware, and performing it in the driver instead.	924738
SR-IOV	Amended the PF\VF counters' wrong data display. For reliable PF/VF counters, it is recommended to use the WinOf 5.40 version on the host, and WinOf 5.30 or higher as the VF version.	965857
	Modified SR-IOV VF driver to use Mellanox hardware comm channel.	853685
mlxtool	Added the ability to show the sensors state for each instance (bus or miniport)	1044293
	Added the ability to modify receive ring parameters on the fly.	1022471
	Exposed Ethernet driver data path statistics via the mlxtool <code>perf-stats</code> command.	1022469
	Added an option to dump all ND connections.	
	Added a command that triggers a "Dump me now" event.	
	Added performing "HCA reset" function.	
	Added the ability to show the difference between actual configured registry keys and values in the registry.	
RDMA/ND	Extended the ND API to allow reporting to RDMA applications when the device is reset, and when it is back to operational mode. The new ND API header is a part of the SDK.	1020585
	Extended the ND API to allow triggering "Dump-me-now" from an ND application for specific QPs.	1022739
	Returning correct return values in case of canceled ND requests.	1022723
	The RoCE Version Interoperability feature is disabled - RDMA communication between nodes, where one node is configured to RoCE v1.5 and the other to RoCE v2, is not supported.	
RSS Monitoring	Added the ability to sample packets which are not sent to the RSS queues via mlxtool.	867203
	Added per-core RSS traffic counters.	867201
ND-IBAL	Removed the option to use or install the ND IBAL provider. In case the ND IBAL provider is installed as part of an upgrade with the full package, it will be removed.	1014850

Table 14 - Change Log History (Sheet 2 of 22)

Category	Description	Reference Number
Troubleshooting	The maximal size of the WPP trace file was increased from 16MB to 100MB.	
Rev. 5.30		
Contains the following versions of components:		
<ul style="list-style-type: none"> • Bus, eth, IPoIB and mux drivers version is 5.35.12978 • The CIM provider version is 5.35.12978 		
Ethernet	Updated driver settings for Virtual Function to receive optimal performance	
	Added RSS (Rx Steering Mode) monitoring counters support	867201
	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.	818135
Tools	vstat tool - Added interface description for each port in the vstat tool.	857255
	Mlxtool - Added support in the following actions: <ul style="list-style-type: none"> • Driver restart • Pulling resource tracker information for VFs 	867812
Bug fixes	See Section 4, “Bug Fixes History”, on page 16	
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 SrioVPortMode 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.	-

Table 14 - Change Log History (Sheet 3 of 22)

Category	Description	Reference Number
Bug Fixes	See Section 4, “Bug Fixes History” , on page 16	-
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 		
VXLAN	Setting the dynamic VXLAN UDP port is now supported for dual-port devices when only a single port is active.	-
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 bit-mask.	-
	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
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 		

Table 14 - Change Log History (Sheet 4 of 22)

Category	Description	Reference Number
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 Ethertypes 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	-
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

Table 14 - Change Log History (Sheet 5 of 22)

Category	Description	Reference Number
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 <code>Set-MlnxDriverCoreSetting</code> 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
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 para-virtualization, 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 <code>mlxtool</code> 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

Table 14 - Change Log History (Sheet 6 of 22)

Category	Description	Reference Number
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
Ethernet	Fixed the adapter name in the <code>vea_man</code> command to case insensitive. 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
	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 14 - Change Log History (Sheet 7 of 22)

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
	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	-
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

Table 14 - Change Log History (Sheet 8 of 22)

Category	Description	Reference Number
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
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 <code>ib_join_mcast()</code> 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 <code>ndkgetremotetokenfrommr()</code> 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 14 - Change Log History (Sheet 9 of 22)

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
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 “Teaming” and “team” respectively.	-
CIM/WMI	Added support to query/set/enable/disable ECN	
	Added support to query DroplessMode state	

Table 14 - Change Log History (Sheet 10 of 22)

Category	Description	Reference Number
	Fixed the issue when using the PowerShell command <code>Get-Mlnx-FirmwareIdentity</code> 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 <code>perf_tune</code>	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: <code>AllowResetOnError</code> 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
IPoIB	Added multiple <code>P_Key</code> 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	-

Table 14 - Change Log History (Sheet 11 of 22)

Category	Description	Reference Number
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 IND2Adapter::QueryAddressList	-
InfiniBand	Added non-default PKey in VM	-
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 		

Table 14 - Change Log History (Sheet 12 of 22)

Category	Description	Reference Number
Generic	Optimized handling of “affinity change” on OID_RECEIVE_FILTER_QUEUE_PARAMETERS	-
	Added the ability to control the number of retries and timeout to check the device health before performing reset	-
Ethernet	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	-
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 part_man 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	-

Table 14 - Change Log History (Sheet 13 of 22)

Category	Description	Reference Number
RoCE	Added ConnectX®-3 Pro support for RoCEv2	-
	Changed the transport name in vstat and ibstat to be RoCE v2.0	-
	Fixed ibstat behavior on devices with RoCE enabled	-
	Fixed releasing of RDMA resources and reacquire them on power down and up.	-
	Fixed RDMA Activity counters which didn't increase for ND traffic	-
ND	Fixed hard-coded limitation of 4 SGEs	-
InfiniBand	Fixed vstat printing of counters for Ethernet ports	-
	Fixed crash when calling <code>ib_join_mcast()</code> with <code>timeout_ms = 0</code>	330740
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 14 - Change Log History (Sheet 14 of 22)

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	-
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	-

Table 14 - Change Log History (Sheet 15 of 22)

Category	Description	Reference Number
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	-
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 14 - Change Log History (Sheet 16 of 22)

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 14 - Change Log History (Sheet 17 of 22)

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 as Interrupt</code> is no longer supported Removed the <code>LsoV1IPv4</code> from the registry/UI Removed from the bus driver configuration the <code>'Non-DMA'</code> 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>*EncapsulatedPacketTaskOffload</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 MLX4ND2 are installed by default Fixed setting the correct SL in UD traffic over RoCE 	-
InfiniBand	<ul style="list-style-type: none"> IPoIB performance improvements Fixed a <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 14 - Change Log History (Sheet 18 of 22)

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 	-
Ethernet	<ul style="list-style-type: none"> • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows’ Perfmon utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive <p>In addition to old values that are not supported anymore.</p> • Made <code>mlx4_bus</code> and Ethernet devices removable • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to <code>mlx4nd2</code> • 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 14 - Change Log History (Sheet 19 of 22)

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 EndPoint with zero attributes in the Receive Queue • Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again • Network Direct: Fixed a failure when creating an EndPoint with zero attributes in the Receive Queue • Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host • Added Transmit Side Scaling (TSS) • Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' Perfmon utility • Added to the MTU size the IP header size (1500 ->1514, 9600->9614). Thus the minimum Jumbo frame size is 614. • Interrupt moderation supports the following profiles: <ul style="list-style-type: none"> • Low Latency • Moderate • Aggressive In addition to old values that are not supported anymore. 	-

Table 14 - Change Log History (Sheet 20 of 22)

Category	Description	Reference Number
Ethernet	<ul style="list-style-type: none"> • Made <code>mlx4_bus</code> and Ethernet devices removable • Network Direct: Added support for NDv2 • Network Direct: Set the default ND provide value to <code>mlx4nd2</code> • 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 <code>MaxRssProcessirs</code> 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 <code>InboundReadLimit</code> and <code>OutboundReadLimit</code> when creating an <code>EndPoint</code> and in <code>Connector::GetConnectionData</code> • Network Direct: Fixed disallowing creation of <code>EndPoint</code> 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 <code>EndPoint</code> limit • Network Direct: Fixed a failure of ND connectivity between VMs on the same host 	
InfiniBand	<ul style="list-style-type: none"> • On rare occasions, depends on the GUID assignment, the IPoIB MAC address can be assigned with a multicast MAC (the least significant bit of the most significant address octet is set to 1). In that case, all of the traffic over the IPoIB I/F is dropped. If you experience this issue, please contact Mellanox support. • Added <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 	

Table 14 - Change Log History (Sheet 21 of 22)

Category	Description	Reference Number
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"> RssMaxProcNumber NumRSSQueues RSSProfile The rest of the keywords are added in all versions of NDIS. This change is based on: http://msdn.microsoft.com/en-us/library/windows/hardware/ff570864(v=vs.85).aspx 	
Ethernet	<ul style="list-style-type: none"> RoCE MTU value is no longer set to 1024 by default. All options stay as they are and can only be chosen if they were selected explicitly in the UI/registry. The current default state is as follows: The value is now derived from the MTU (or MaxFramSize, or Jumbo Packets value) and they are all aliases for the same value). The value is aligned to 256,512,1024,2048 in a way that it will be rounded down to the nearest power of two of the ETH MTU. 	
InfiniBand	<ul style="list-style-type: none"> Added ibdiagnet utility support 	
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 14 - Change Log History (Sheet 22 of 22)

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_I-PoIB • for ConnectX-3: MLX4\ConnectX-3_Eth and IBA\ConnectX-3_I-PoIB • 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 processor 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 	
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 	

6 API Change Log History

Table 15 - API Change Log History

Release	Name	Description
5.40.50000	ND IBAL provider	Disabled ND IBAL provider.
	Extended ND API	Allows reporting to RDMA applications when the device is reset, and when it is back to operational mode. For more information see the User Manual.
4.95.50000	ND extension for Resource pools	Please refer to MLNX_VPI_WinOF_User_Manual_v4.95
4.80.50000	RDMA_TRANSPORT_RDMAOE_1	It is an alias to: RDMA_TRANSPORT_RDMAOE
	RDMA_TRANSPORT_RDMAOE_1_25	Added enumerated values
	is_roce(), is_mac_based_roce(), is_ip_based_roce(), is_rroce_or_ip_based_roce()	Added new functions
	struct ib_wc_t	p_next was replaced with an anonymous union which contains two fields: p_next and qp_context
4.70	ib_get_port_spl_qp()	Added a new function
	ib_get_mad_inner()	Changed API (one more input parameter was added)
	ib_get_mad()	Changed API (one more input parameter was added)
	VERBS_MINOR_VER	Increased its value, 0x000a -> 0x000c
	UNBOUND_PORT_NUM	Added a new macro
4.60	IB_MOD_QP_CHANGE_COUNTER_INDEX	Added a new macro
	struct ib_qp_mod_t	Added the field state.rtr.counter_index
4.55	RDMA_TRANSPORT_RRDMAOE_1_5	Added enumerated values
	RDMA_TRANSPORT_RRDMAOE_2_0	Added enumerated values
4.50	is_rroce(), is_xroce()	Added new functions
	IB_AC_SNIFFER	Added a new macro
	struct ib_qp_mod_t	Added the field state.init.flags

Table 15 - API Change Log History

Release	Name	Description
4.40	VERBS_MINOR_VER	Increased its value, 0x0009 -> 0x000a
	enum eth_link_speeds	Added enumerated values
	struct ib_port_attr_t	<ul style="list-style-type: none"> The mtu field was separated into two fields: <ul style="list-style-type: none"> max_mtu (maximum MTU supported by the port) active_mtu (actual MTU which the port is configured with) Added the eth_link_speed field
	WR_SEND_INV	Added enumerated values
	struct ib_send_wr_t	The type of invalidate_rkey was changed from net32_t -> ib_net32_t
	IB_SEND_OPT_SKIP_DOORBELL	Added the send Write flag