

Connect. Accelerate. Outperform.™

Mellanox WinOF VPI Release Notes

Rev 4.90.50000

www.mellanox.com

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 NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies 350 Oakmead Parkway Suite 100 Sunnyvale, CA 94085 U.S.A. www.mellanox.com Tel: (408) 970-3400 Fax: (408) 970-3403 Mellanox Technologies, Ltd. Hakidma 26 Ofer Industrial Park Yokneam 2069200 Israel www.mellanox.com Tel: +972 (0)74 723 7200 Fax: +972 (0)4 959 3245

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

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

CyPUTM, ExtendXTM, FabricITTM, FPGADirectTM, HPC-XTM, Mellanox CareTM, Mellanox CloudXTM, Mellanox Open EthernetTM, Mellanox PeerDirectTM, Mellanox Virtual Modular SwitchTM, MetroDXTM, NVMeDirectTM, StPUTM, Switch-IBTM, Unbreakable-LinkTM are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Chapter 1	Ove	rview	. 3
	1.1	Main Features in this Release	. 3
	1.2	Beta Features	. 3
	1.3	Unsupported Functionality/Features	. 3
	1.4	WinOF VPI Package Contents	. 4
	1.5	Supported Operating System Versions	. 4
	1.6	Supported Network Adapter Cards	. 5
	1.7	RoCE Modes Matrix	. 5
Chapter 2	Cha	nges and Fixes in Rev 4.90.50000 from v4.80.50000	. 6
	2.1	API Changes in WinOF Rev 4.90.50000	. 8
Chapter 3	Kno	own Issues	. 9
	3.1	Generic Issues	. 9
	3.2	InfiniBand Issues	11
	3.3	Ethernet Issues	14
	3.4	Quality of Service Issues	16
	3.5	RoCE NVGRE Issues	16
	3.6	Performance Issues	18
	3.7	Hyper-V Issues	18
	3.8	SR-IOV Issues	21
	3.9	Installation / Upgrade Issues	23
	3.10	Utilities Issues	27
	3.11	CIM/WMI Issues	27
	3.12	ND Issues	28
Chapter 4	Cha	nge Log History	29
Chapter 5	API	Change Log History	50

1

List of Tables

Beta Features
Supported Operating Systems
Supported Network Adapter Cards
Firmware Versions
RoCE Modes Matrix
Changes and Fixes in Rev 4.90.50000
API Changes in WinOF Rev 4.90.50000
Generic Issues
InfiniBand Issues
Ethernet Issues
Quality of Service Issues
RoCE NVGRE Issues
Performance Issues
Hyper-V Issues
SR-IOV Issues
Installation / Upgrade Issues
Utilities Issues
CIM/WMI Issues
ND Issues
Change Log History
API Change Log History

1

Release Update History

Release	Date	Description
Rev 4.90.50000	June 15, 2015	Added one known issue to Table 10 - "Ethernet Issues," on page 14
	April 27, 2015	Clarified some fixes/changes in Section 2, "Changes and Fixes in Rev 4.90.50000 from v4.80.50000", on page 6
	January 2015	Initial Release

1 Overview

These are the release notes for Mellanox WinOF Rev 4.90.50000 VPI drivers.

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

1.1 Main Features in this Release

WinOF Rev 4.90.50000 provides the following new features:

• UDP Source port:

Routable RoCE UDP Source port value is now taken from the ND layer

• ETS with DSCP:

Added the ability to control ETS using DSCP values

• IPoIB failover:

Allows the creation of IPoIB interface that binds several physical IPoIB interfaces for fault tolerance

• Multiple PKeys:

Added IPoIB support for use of multiple interfaces with different PKeys on the same port

1.2 Beta Features

Table 1 - Beta Featu	ıres
----------------------	------

Category	Description	
WinVerbs	WinVerbs is currently at beta level.	
ibdump	ibdump is currently at beta level.	
IPoIB	 IPv6 support of IPoIB in an SR-IOV guest OS over KVM is at beta level. IPoIB teaming support is at beta level and it is supported only on native machine (and not in HyperV or SR-IOV). 	

1.3 Unsupported Functionality/Features

The following are the unsupported functionalities/features in WinOF Rev 4.90.50000:

- ND over WinVerbs provider
- SRP
- IPv6 over IPoIB Team ports
- VMQ over IPoIB team ports
- Configure IPoIB team through PowerShell
- ConnectX®-2 adapter cards

1.4 WinOF VPI Package Contents

The Mellanox WinOF Rev 4.90.50000 for Windows package contains the following components:

- Core and ULPs:
 - IB HCA low-level drivers (mlx4)
 - IB Access Layer (IBAL)
 - Ethernet driver (ETH)
 - IP over InfiniBand (IPoIB)
 - NetworkDirect (ND)
- Mellanox LBFO driver for Windows Server 2008 R2
- Utilities:
 - Low level performance tools
- CIM, PowerShell, and WMI support

1.5 Supported Operating System Versions

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

Virtualization Mode	Supported Host OS	Supported Guest OS
Native (no-virtualiza- tion)	Windows Server 2008 R2 (64 bit only)	N/A
	Windows Server 2012 (64 bit only)	N/A
	Windows Server 2012 R2 (64 bit only)	N/A
Hyper-V (non-SR- IOV)	Windows Server 2008 R2 (64 bit only)	• Windows Server 2008 R2 (64 bit only)
	Windows Server 2012 64 bit only)	 Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only)
	Windows Server 2012 R2 (64 bit only)	 Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)
SR-IOV Ethernet	Windows Server 2012 R2 (64 bit only)	 Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)
SR-IOV InfiniBand	KVM with MLNX_OFED 2.4-x.x.x	 Windows Server 2008 R2 (64 bit only) Windows Server 2012 (64 bit only) Windows Server 2012 R2 (64 bit only)

Table 2 - Supported Operating Systems

1.6 Supported Network Adapter Cards

Mellanox WinOF Rev 4.90.50000 supports the following Mellanox network adapter cards:

Table 3 - Supported Network Adapter Cards

NICs	Supported Protocol	Supported Link Speed
ConnectX®-3 Pro	InfiniBand (IB)	10, 40 and 56 Gb/s
	Ethernet	10, 40 and 56 Gb/s
ConnectX®-3	InfiniBand (IB)	10, 40 and 56 Gb/s
	Ethernet	10, 40 and 56 Gb/s

1.6.1 Firmware Versions

Mellanox WinOF Rev Rev 4.90.50000 provides the following firmware for Mellanox NICs:

Table 4 - Firmware Versions

NICs	Recommended Firmware Rev.	Additional Firmware Rev. Supported
ConnectX®-3 Pro / ConnectX®-3 Pro EN	Rev 2.33.5100	Rev 2.32.5100
ConnectX®-3 / ConnectX®-3 EN	Rev 2.33.5100	Rev 2.32.5100

1.7 RoCE Modes Matrix

The following is RoCE modes matrix:

Table 5 - 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)	4.70 (Requires additional configuration)
Inbox Windows Server 2012 / Inbox Windows Server 2012 R2	Supported (Default)	Not supported	Not supported

2 Changes and Fixes in Rev 4.90.50000 from v4.80.50000



This package version is Rev 4.90.50000. The package contains the following versions of components.

• Bus, eth, IPoIB and mux drivers version is 4.90.10714

• The CIM provider version is 4.90.10714

Table 6 - Changes and Fixes in Rev 4.90.50000

Category	Description
Generic	 Fixed driver instability when handling many RDMA connection requests in parallel Added to MLNX_System_Snapshot Mellanox specific counters and data from Get- Mlnx* Cmdlets
Resiliency	Reset Flow improvements:Resolved race condition when reset is initiated by more than one sourceReset initiated on one port does not cause reset of the other port
Ethernet	 ETS is now configurable through DSCP values. For further details, please refer to WinOF User Manual, "Differentiated Services Code Point (DSCP)" section Fixed the issue of when creating a Virtual Ethernet Adapter interface and removing it immediately a Blue Screen may appear Fixed duplicated values of Receive Completion Method in Advanced Properties driver dialog on Windows Server 2012 R2 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
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 Fixed propagation of error code when ib_join_mcast() fails Fixed connectivity problems when using PKeys from the same partition with different membership types 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
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 hanged ndkgetremotetokenfrommr() to return value in network byte order

Table 6 - Changes and Fixes in Rev 4.90.50000

Category Description		
Performance	 Fixed the UI crash when working with a single port 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 	
IPoIB	 Fixed stability issues Fixed displaying of IPoIB default turning option Fixed displaying of IPoIB default turning option Fixed temporary network connectivity issues while migrating VMs or modifying VMQ configuration for VMs that uses IPoIB with VMQ Fixed the part_man tool to use the actual default p_key instead of 0xffff Fixed NIC reset when attaching to a multicast group fails 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 Added sending of gratuitous ARP in IPoIB interface when the MAC address is changed Reduced the multiple number of path record queries to one when old query information exists Improved completion memory access speed Changed default VMQ/VPort affinity to use first RSS CPU Multiple PKey support is now at GA level. The part_man tool allows the creation of up to 64 vIPoIB interfaces (32 per port) 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 Fixed loading of old driver after driver upgrade that requires system reboot to complete the process Fixed RoCE disable by default after installation of WinOF in Windows 8.1 Client 	
ND	 Fixed seg fault when executing ND application with no device installed or when a wrong device identifier is used Fixed wrong reported value of supported number of SGE in 32 bit DLLs Increased the number of supported SGEs in 32 bit DLLs to 2 	
NVGRE	Fixed restoration of NVGRE configuration after NIC reset	
Changes in UI	Replaced the terms "LBFO" and "Bundle" with "Teaming" and "team" respectively.	
CIM/WMI	Fixed failure when calling Get-MlnxFirmwareIdentity if not all devices are up	
	Fixed the issue when using the PowerShell command Get-MlnxFirmwa- reIdentity on a system with multiple NICs/HCAs while one of the devices is disabled and the command fails	

2.1 API Changes in WinOF Rev 4.90.50000

The following are the API changes in WinOF Rev 4.90.50000

Table 7 - API Changes in WinOF Rev 4.90.50000

Name	Description
struct ib_mad_element_t	Added the field rroce_udp_s_port
struct ib_av_attr_t	Added the field rroce_udp_s_port
VERBS_MINOR_VER	Increased its value, 0x00c -> 000d

3 Known Issues

3.1 Generic Issues

Table 8 - Generic Issues

Issue	Workaround
Pinning all the physical memory (used by RDMA operations, such as register memory, pin user memory) on the machine, on Operating Sys- tems prior to Windows Server 2012, may cause the machine to hang.	Avoid pining the whole machine memory in those Operating Systems.
When running applications that use ND or libibumad (such as OpenSM) the system might get to an unstable state when trying to shutdown/ restart/hibernate it.	Close all applications that use ND or libibumad before performing shutdown/restart/hibernate.
Activating NC-SI in WinOF v4.90.10541 may cause driver's loading failure when using an older firmware version than 2.30.8000.	Do not enable NC-SI in machines that WinOF v4.90.10541 is installed in.
The maximum values returned by the ib_que- ry_ca() function (for example: max_qp, max- _mr) 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 configu- ration, the amount of host memory, user permis- sions, and the amount of resources already in use by other users/processes.	-
Running Ntttcp without the "-a x " flag (X >1) in a NIC configured with 10GbE, may cause low bandwidth in TCP single stream.	Run Ntttcp with "-a 8" for best performance
Active links disappear after changing the cable connectivity from Ethernet to InfiniBand or vice versa.	Disable and enable the mlx4_bus interface from the Device Manager.

Table 8 - Generic Issues

Issue	Workaround
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.	Please use the GUID value returned by the fab- ric/driver utilities (not 0xfffff).
The UI displays the network interface's state incorrectly even after disabling and re-enabling the "disconnected" network interfaces.	 To see the correct state and link speed, perform one of the following: Run Get-netadapter Powershell cmdlet or Right click on that interface from "Network Connections" and click on status
WoL is not supported on Windows Server 2008 R2.	-
iSCSI boot over Windows Server 2008 R2 is not supported.	-
Clearing the Mellanox device counters through perfmon does not always work.	Restart the driver
When running in SR-IOV mode and Hyper-V, the same driver version should be used on both the Hypervisor and the virtual machine.	-
Working with user space RDMA resources (i.e. ND) in parallel with hardware reset may cause unexpected behavior.	-
In Windows Server 2008 R2 and Windows 7 Cli- ent, ConnectX-3 Pro is identified as ConnectX-3. This does not affect anything.	-
On rare occasions, as a result of port configura- tion 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.	-
The drivers' functionality is limited up to 128 cores.	-

Table 8 - Generic Issues

Issue	Workaround
According to the hardware architecture, Mella-	Any customer programs or scripts that look into
nox NIC devices are presented in the registry as	the NIC registry values should not assume it is a
virtual devices and not as physical devices.	physical device.

3.2 InfiniBand Issues

Table 9 - InfiniBand Issues

Issue	Workaround
IPv6 traffic between Hyper-V hosts over IPoIB v-Switch may experience traffic loss.	-
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.	Remove the virtual port before the upgrade and re-add it after the upgrade.
InfiniBand application that was compiled with an SDK version earlier than WinOF v4.90 is not binary compatible.	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.
WinOF 4.40 and above IPoIB driver is not IPv6 compatible with earlier driver versions, includ- ing Window 8 Inbox driver. If WinOF 4.50 IPoIB node receives an icmpv6 message from the legacy IPoIB node, the follow- ing event will appear in the event log: "IPoIB driver detected a possible incom- patibility with Windows 8 inbox IPv6 support due to which there is no IPv6 connectivity".	 To enable compatibility mode, add: Win8InboxCompatibilityMode REG_SZ registry key with the value of 1 in the IPoIB interface registry. 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
Without separate ports for each stream, WinSock multiplexes every packet to every subscriber socket and then filters it out.	Use different UDP ports to get higher perfor- mance when using multicast packets.
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	-
The "Packets Received Discarded" and "Packets Received Errors" counter may display wrong results.	-

Table 9 - InfiniBand Issues

Issue	Workaround
Connection failure on ND tests while machine A have IBAL provider and machine B have MLX-4ND provider.	-
Hibernate and Sleep are not functional when user-space is using its resources.	-
IPoIB does not support:MAC address changeQoS (packet priority)Connected Mode	-
In an interoperability environment that has both Linux and Windows OSs, the MTU value must be the same, otherwise packets larger than the minimum will not go through. The default MTU for Linux is 2K and for Windows is 4K.	-
OpenSM does not run as a service during instal- lation since the SM requires the GUID parameter to decide which port to work on. Setting it on setup causes it to work only on the first port and not the others.	To run OpenSM as a service, assuming the pack- age was installed in the default path, use: sc create opensm binPath= "c:\Program Files\Mellanox\ MLNX_VPI\IB\Tools\opensm.exe" To start the service run: sc start opensm
Tools issues:ibportstate does not work on RoCE portsibdiagpath may crash on Hyper-V machines	-
If an application which uses InfiniBand runs while the driver is being restarted, a bluescreen or an NMI may occur.	Stop all InfiniBand applications including OpenSM upon driver restart.
If OpenSM is up during driver restart on the same machine, it might stop working, and an error message that reads "Another OpenSM is running" will pop up.	To resume operation, stop OpenSM and restart the driver while OpenSM is down.
Sometimes during traffic, the latency of the IPoIB can get higher (this behavior is inconsistent).	Set the following registry in the IPoIB network interface: RecvIntModTime=0 Please note, this action increases the CPU utiliza- tion.
No communication between the physical inter- face and a VM that uses vSwitch created over virtual IPoIB (Pkey), and vice versa.	-

Table 9 - InfiniBand Issues

Issue	Workaround
Burning a Firmware image with a "mtu_cap" value other than the default causes the driver load to fail.	Do not change the "mtu_cap" value
Calling ib_query_cq() on a CQ which was cre- ated with 1 entry indicates that there are 0 entries in the CQ.	-
IPoIB ports report the configured MTU instead of the effective MTU. These MTUs differ when the corresponding InfiniBand partition is config- ured with a smaller MTU than the port's config- ured MTU. In such case, there may be communication failures and/or CPU work over- head on the communicating nodes (even if the peer node is configured correctly).	Make sure the configured adapter MTU (the "Jumbo packet" property in the adapter's advanced settings) is consistent with the corre- sponding 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 net- work switches) on the InfiniBand subnet needs to be changed. Note that the adapter's MTU can be configured through PowerShell.
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 func- tional.	-
DHCP messages that IPoIB send are not fully spec complaint. The spec specifies that the 'chaddr' field must be zeroed, but WinOF IPoIB does not guarantee it.	-

3.3 Ethernet Issues

Table 10 - Ethernet Issues

Issue	Workaround
Disabling the "Priority & VLAN tag" in the UI which VLANID is configured, may result in sending packets with the configured VLANID.	Remove the VLANID before disabling the "Pri- ority & VLAN tag".
When working with LBFO, the teamed interface disappears after machine reboot. The issue applies to Windows Server 2012 R2	Delete the existing teamed interface and create a new one.
WakeOnMagicPacket registry key is not added to the registry although WoL is supported by the driver and by the NIC.	-
When the ports of the device are configured as Ethernet only, ibstat/vstat may display wrong information.	-
High multicast drop rate on multicast storming.	Use "Multicast traffic" tuning option under the performance tab. For further information, please refer to section "Tunable Performance Parameters" in the User Manual.
When there is a stress in TCP connection estab- lishments, some of those connections may fail.	 Increase the Ring queue sizes: ReceiveBuffers - controls the receive ring size TransmitBuffers - controls the transmit ring size
 The DCB component specifies a default traffic classification that is applied to all egress packets that do not match other classification conditions. In this case, the network adapter assigns the IEEE 802.1p priority level that is associated with the default classification to these egress packets. The default traffic classification has the following attributes: 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_CLASSIFICA-TION_ELEMENT structures. 	-

Table 10 - Ethernet Issues

Issue	Workaround
Disabling the ethernet adapter with more than 25 VLANs configured over windows server 2008 R2 \ windows server 7, may result with a non- responding server.	Reduce the number of configured VLANs prior to disabling the ethernet adapter.
Virtual Ethernet Interfaces created by vea_man are not tuned by the automatic performance tun- ing script.	For optimal performance need to follow the per- formance tuning guide and apply relevant changes to the VEA interface
In Windows Server 2008 R2, devices created by the Teaming driver do not show the correct OEM branding in the Device Manager.	-
Wake on Lan (WoL) cannot be disabled on NICs which supports it.	-
Raising the Receive Buffers default value (512) could cause the system to crash.	-

3.4 Quality of Service Issues

Table 11 - Quality of Service Issues

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

3.5 RoCE NVGRE Issues

Table 12 - RoCE NVGRE Issues

Issue	Workaround
 RoCE does not support: Traffic cannot go through the router. It works in the same subnet only Multicast traffic VLAN Layer 3 feature 	-
 In machines with heterogeneous NICs: a NIC which supports RoCE v2, and a NIC which does not support RoCE v2 the following issus might raise: ConnectX®-3 Pro is loaded with the transport type RoCE v2 ConnectX®-3 Pro is loaded with the transport type RoCE v2 	-
When NVGRE off-load is enabled, the GRE traf- fic 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.	-
Using different versions of RoCE in your cluster is not supported.	Use the same RoCE version in all the cluster in the Ethernet ports.
RDMA Activity counters do not count during NetworkDirect RoCE traffic.	-
GRE traffic steering by inner MAC and by outer MAC simultaneously is currently not supported.	Configure steering or by inner MAC, or by outer MAC.

Table 12 - RoCE NVGRE Issues

Issue	Workaround
If VMQ set filter requests are accepted without a GRE flag (i.e. requested steering by outer MAC), the GRE packets do not reach that VMQ.	Set the bus driver registry keyAcceptGREby- OuterMAC_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 fil- ter 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

3.6 Performance Issues

Table 13 - Performance Issues

Issue	Workaround
When using WinOF 4.40 or above, low through- put will be seen on 40GbE adapters when QoS is enabled.	<pre>Disable QoS when it is not in use: Open a PowerShell prompt. Run: Disable-NetAdapterQos -name <inter- face="" name=""> where <interface name=""> is e.g. "Ethernet 1"</interface></inter-></pre>
perf_tuning is supported only when one of the two NUMA nodes are in use.	-
Running performance benchmarks for a short period of time (< 1 sec) may provide bad latency in IPoIB and Ethernet.	Set "Rx Interrupt Moderation Profile" and "Tx Interrupt Moderation Profile", to "Low Latency" to avoid bad latency. Note: This may increase CPU utilization.
The driver uses optimal interrupt moderation val- ues for 10 GbE SR-IOV VF scenario. For other scenarios, the optimal values yet to be found.	-
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.	Reduce the counters sample rate.

3.7 Hyper-V Issues

Table 14 - Hyper-V Issues

Issue	Workaround
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.	-

Table 14 - Hyper-V Issues

Issue	Workaround
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 pack- ets will be passed from NDIS down to the driver and they will be dropped by it.	Reattach the vSwitch to sync with the value set in the ETH\IPoIB device.
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.	-
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	-
Hyper-V is at low bandwidth on LBFO vSwitch, Windows Server 2012	-
In Ethernet to achieve better iperf TCP perfor- mance 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.	-
After disabling and enabling a port on a guest, a ping to it may be renewed after a minute. The ARP requests sent by Windows are less frequent as the time passes. If the guest port was down for a while, it could take time until Windows decides to send another ARP request to it.	-
When VMQ is enabled after reset, the driver loads all the VMQs that existed before the reset. However, it is not guaranteed that each VMQ will receive the same QP number it had before the reset. This can cause some delay as a result of resetting before connectivity is reestablished. The delay is caused by the time it takes for the ARP table to update after initiating the Gratuitous ARP.	-

Table 14 - Hyper-V Issues

Issue	Workaround
The IPoIB non-VMQ mode is supported only when the VMQ is enabled according to the regis- try values.	 To use the non-VMQ mode for a VM, change its settings as follow: Press "Settings" on the VM Go to Network Adapter -> Hardware Acceleration Un-check the "Enable virtual machine queue"

3.8 SR-IOV Issues

Table 15 - SR-IOV Issues

Issue	Workaround
Working with Windows guest OS over non-Win- dows SR-IOV hypervisor may result with higher latency compared to Windows Hypervisor.	-
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.	Verify that the FW INI contains the parameter vf_subsystem_id with the OEM encoding in the [HCA] section.
An SR-IOV Virtual Machine is loaded in Infini- Band mode when no OpenSM is available in the subnet. The following event appears in the event log: "detected a null port GUID for port <num- BER>. 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."</num- 	Start the OpenSM and restart the driver.
For InfiniBand SR-IOV guest, OpenSM Assigned GUIDs are not supported and may cause unexpected behavior	Work only with Administrator assigned GUIDs.
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.	Disable the Para-Virtualization adapter.
In SR-IOV mode over Hyper-V, all ports are in pure Ethernet mode and RDMA is not supported on either port.	-
In SR-IOV mode, enabling SR-IOV in the BIOS may change the interfaces names. If any VSwitch is bounded to an interface whose name was changed, there will not be any way to perform any operation on them.	Unbond all Vswitches from the NIC's interfaces before enabling SR-IOV in the BIOS.
In SR-IOV, 40Gbps interfaces of vSwitch inter- face on HyperV, and the VMNIC are identified as 10Gbps. Despite the incorrect speed identifi- cation, they still achieve 40Gbps performance.	-
Device name in the Device Manager does not show the correct OEM branding for SR-IOV Vir- tual Function devices without the updated firm- ware.	-

Rev 4.90.50000

3.9 Installation / Upgrade Issues

Table 16 - Installation / Upgrade Issues

Issue	Workaround
Upgrading the driver while the UI is opened with the "ConnectX NIC device" may cause the installation process to never end.	Close the UI before driver upgrade.
Rebooting the machine while uninstalling WinOF may result in installation failure.	Delete Mellanox components from HKEY_LO- CAL_MACHINE\SOFTWARE\Microsoft\Win- dows\CurrentVersion\DIFxApp\Components. The Mellanox components are mlx4eth63, ipoib6x and mlx4_bus
Canceling the installation process may leave the bus driver in a disable state. The driver appears in a yellow bang containing the following error message: "Windows cannot start this hardware device because its configuration infor- mation (in the registry) is incomplete or damaged. (Code 19)".	Scan for new hardware and reboot the machine.
Downgrade is not supported.	Uninstall the current version and install the older one.
Configuration is not restored when replacing a ConnectX®-3 NIC with a ConnectX®-3 Pro NIC located on the same PCI slot	Clean the old network adapter configuration prior to upgrade.
Configuration can be restored only in Windows Server 2012 and above	-
IPv6 configuration restore is not supported	-
Upon upgrade, the following Registry Key val- ues will be overwritten with the following: *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	

Table 16 - Installation / Upgrade Issues

Issue	Workaround
Upon upgrade the following Ethernet Registry Keys will be deleted: • SendCompletionMethod • UseRSSForRawIP • UseRSSForUDP	-
 Upon upgrade the SendCompletionMethod IPoIB Registry Key value will be modified as follow: SendCompletionMethod = 0 	-
 Upon upgrade the following IPoIB Registry Keys will be deleted: UseRSSForRawIP UseRSSForUDP 	-
Uninstalling the driver on Windows Server 2008 R2 with LBFO configuration results in the appearance of a pop-up window requesting to close several running applications.	Choose "Do not close applications". This action allows the uninstallation of the driver. A Reboot may be required. Rebooting the server before uninstalling the driver when LBFO is configured will eliminate this pop-up completely.
Running a downgrade in silent mode is not sup- ported. Upon downgrade the return code will always be 0.	-
Uninstalling the driver after upgrade won't remove the directory %ProgramFiles%\Mella- nox	-
Uninstalling the driver when multiple of VLANs are configured never ends.	Remove the VLANs before uninstallation.
The installation process does not close any appli- cations running in the background, and may cause a BSOD as a result of a stuck cmd.	It is recommended to close all running applica- tions prior to upgrading the driver.
<pre>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\DMI1 51A.tmp.log.xml C:\Program Files\Mellanox\MLNX- _VPI\ETH\mlx4eth63.inf</user></pre>	Reboot the machine and reinstall

Table 16 -	Installation	/ Upgrade i	lssues
------------	--------------	-------------	--------

Issue	Workaround
Installation/upgrade fails due to failure to stop the WMI service, and the following text is printed in the installation log: "CustomAction StopWMIService returned actual error code 1603"	Kill the WMIPrvSE.exe tasks in the task man- ager and reinstall.
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 com- pleted. 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?TOP- IC ID=39253	<pre>Run the following script on the server to opti- mizes loading PowerShell DLLs: \$Env:PATH = [Runtime.InteropSer- vices.RuntimeEnvironment]::GetRuntime- Directory() [AppDomain]::CurrentDomain.GetAssem- blies() % { \$pt = \$Location if (! \$pt) {continue} if (\$cn++) {''} \$na = Split-Path -Leaf \$pt Write-Host -ForegroundColor Yellow "NGENing \$na" ngen install \$pt }</pre>
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 con- figured following the upgrade.	Reconfigure the RoCE Mode setting manually.
Uninstall from the Device Manager is currently not supported	-
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 auto- matically when the new driver is uninstalled.	-
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 installa- tion.	Upgrade the firmware manually.
Driver installation requires deletion of the mlx- 4_bus.sys file in Windows Server 2008 R2 and WLH OSs when using the PXE package.	Delete the mlx4_bus.sys file and reboot the machine to install the driver.

Table 16 - Installation / Upgrade Issues

Issue	Workaround
If there are disabled network interfaces or a dis- abled Mellanox bus driver, they will be enabled after the WinOF upgrade.	-
The upgrade from WinOF v4.58 to WinOF v4.90 saved only the configuration of the Ethernet interface, and did not save the configuration of the IPoIB interface.	

3.10 Utilities Issues

Table 17 - Utilities Issues

Issue	Workaround
ibdump may encounter packet drops upon a burst of more than 4096 (or 2 [^] max-burst) pack- ets.	-
Packets loss is not reported by ibdump.	-
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.	Verify Ethernet flow control is enable to ensure a lossless link
 Pcap file issues for RoCE IP Based: 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'. 	-
Sniffing over IB ports is currently not supported	-
Creating virtual adapters over IPoIB using part_man is currently not supported in all OEM adapter cards.	-
When the tracer tool traces many events, it may consume a large amount of memory (up to sev- eral GB RAM).	Use the button to set maximum number of the displayed lines at a maximum buffer size.

3.11 CIM/WMI Issues

Table 18 - CIM/WMI Issues

Issue	Workaround
Running Microsoft CIM cmdlets operations and their derived classes on classes MLNX_NetAdapter- SettingData and MLNX_NetAdapterRoceSet- tingData is not supported. Calling those commands may cause the debugger, if connected to the machine, to assert.	Use DriverCoreSettings instead.
For PCI Gen3, PcieLinkSpeed is reported as "Unknown" when running Get-NetAdapter- HardwareInfo Powershell cmdlet	-

Table 18 - CIM/WMI Issues

Issue	Workaround
WMI does not work due to lack of permissions.	Change the execution policy. Run: Set-ExecutionPolicy AllSigned
The information that is printed in the cmdlets get-netadaptersriov and Get-MlnxPCIDeviceSet-ting is inconsistent.	Use only the cmdlet Get-MlnxPCIDeviceSet- ting.

3.12 ND Issues

Table 19 - ND Issues

Issue	Workaround
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 oper- ations.	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.
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.	-
Changing the default ND providers may cause random errors, such as: 0xC0000238 (NT_STATUS_ADDRESS_AL- READY_ASSOCIATED) on Connect() or with 0xC0000236 (NT_STATUS_CONNECTION_RE- FUSED) on Accept(). These errors can be safely ignored.	-

4 Change Log History

Release	Category	Description	Notes
4.80.50000	Installa- tion/ Upgrade	 Added check for administrator privileges during installation Added support for installation in silent mode without execution of perf_tune Fixed installation stuck when Remote Desk- top Session Host Windows Installer RDS compatibility is enabled 	 This package version is 4.80.50000. The package contains the following versions of components. Bus, eth, IPoIB and mux drivers version is 4.80.10388 The CIM provider version is 4.80.10388
	Generic	 Changed Reset Flow (+SR-IOV)-enabled only if no user space application is running and depends on the registry key: AllowRese- tOnError setting) Changed the number of supported QPs in a multicast group from hard coded value to firmware capabilities dependent Fixed driver load failure in machines with 1 TB memory and above Fixed memory leak on the Virtual Machine in SR-IOV when resetting the Virtual Machine of associated VFs 	
	IPoIB	 Added multiple P_Key support (beta level) Added IPoIB SR-IOV over KVM and ESX Hypervisors (for both full and partial mem- bership) Added support for LID change event Added enhancements in part_man 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 	

Release	Category	Description	Notes
	Ethernet	 Added RSS in UDP (enabled by default) Added 56 GbE (Please refer to the Infiniband Switch User Guide for further details) Changed DSCB configuration of the second s	
		Changed DSCP configuration to be per port instead of global	
		Network Direct: Fixed race in NDK between handling of incoming connection and destruc- tion of a listener	
		Network Direct: Fixed race between NDK object creation and usage	
		Improved TCB (Transmission Control Block) management on send	
		• Improved transmit and receive in multi stream scenarios	
		Enabled hardware checksum offload for non TCP/UDP traffic with ConnectX®-3 Pro	
		Improved stability when handling OIDs during driver reset	
		• Fixed performance tuning for 1GbE link	
		• Fixed possible reset of driver during migra- tion of large number of VMs at the same time	
		Fixed stability issues	
	RoCE	Added RoCE IP based	
	ND	Fixed wrong return value in IND2- Adapter::QueryAddressList	
	InfiniBand	Added non-default PKey in VM	
	Perfor- mance	Optimized interrupt moderation values in SR- IOV VF mode for IPoIB	
		• Improved perf_tuning detection for the first port	
		• Improved performance in packet forwarding scenarios	
		• Decreased dropped packets rate for Ethernet significantly	
		• Changed default perf_tuning scenario to be "Balanced configuration"	
		• Various performance improvements	
	WMI/CIM	• Added ability to read active RoCE configura- tion from hardware	
		• Added support for RoCE IP Based	

Table 20 -	Change	Log History
------------	--------	-------------

Release	Category	Description	Notes
4.70.50050	IPoIB	• Fixed SM fail-over causing the driver to hang	 This package version is 4.70.50050. The package contains the following versions of components. Bus, eth, IPoIB and mux drivers version is 4.70.10143 The CIM provider version is 4.70.10143
4.70.50040	Generic Ethernet	 Optimized handling of "affinity change" on OID_RECEIVE_FILTER_QUEUE_PA- RAMETERS Added the ability to control the number of retries and timeout to check the device health before performing reset Fixed missing pause response by sender when using DSCP/untag priority tag mode with ETS enabled 	 This package version is 4.70.50040. The package contains the following versions of components. Bus, eth, IPoIB and mux drivers version is 4.70.10141 The CIM provider version is 4.70.10141

Release	Category	Description	Notes
4.70.50000	Installa- tion/ Upgrade	 Fixed removal of virtual IPoIB ports in unin- stallation 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 	 This package version is 4.70.50000. The package contains mixed versions of components. Bus, eth, IPoIB and mux drivers version is 4.70.10126. The CIM provider ver-
	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 	sion is 4.70.10130.
		IPoIB	 Added support for IPoIB SR-IOV Virtual Function (VF) over KVM Hypervisor (Beta level) Added support for non-default pkey, as que- ried from OpenSM, on IPoIB SR-IOV VF over KVM. Added IPoIB QoS proprietary counters, diag- nostics and traffic for monitoring, using Win- dows' 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 Win- dows Server 2012 R2)* Added Virtual Ethernet Adapter support which enables using SMB Direct and HyperV (VMQ and NVGRE (over ConnectX®-3 Pro)) on the same port** Added lossless TCP buffer management when no receive WQE are available 	 * Requires firmware v2.30.8000 and above ** Requires firmware v2.31.5050 and above.

Table 20 -	Change	Log History	'
------------	--------	-------------	---

Release	Category	Description	Notes
	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 reac- quire 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 ib_join_m- cast() with timeout_ms = 0 	
	Perfor- mance	• Improved perf_tuning setting in single CPU machines to avoid TX collision	
4.61	Installa- tion/ 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 	WinOF VPI version 4.61 was released as an interme- diate release.
	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 Mella- nox devices to be listed in "Devices and Printers" and had "Safe Removal" UI 	
	Perfor- mance	 Added support OF IPv6 to all nd_*_* tests Enabled optimal interrupt modera- tion values in SR-IOV VF mode Stopped using NdisQueryNetBuffer- PhysicalCount to improve CPU uti- lization 	

Release	Category	Description	Notes
	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 	
	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 	

Release	Category	Description	Notes
4.60.17738	IPoIB	 Fixed using CQ after VMQ is closed Fixed bad completion of VMQ QP that was caused by malformed WR 	This package version is 4.60.17738. The package contains the following ver-
	Ethernet	• LBFO: Fixed the team's MAC address uniqueness in the subnet of the team in Windows Server 2008 R2	 sions of components: Bus and eth driver ver- sion 4.60.17718. The CIM provider ver- sion is 4.60.17718. The mux driver version is 4.60.17729. The IPoIB driver ver- sion is 4.60.17736.

Release	Category	Description	Notes
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	This package version is 4.60.17718. The package contains the following ver-
	Installa- tion/ 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 Con-nectX-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 	 sions of components: Bus, eth, IPoIB and mux drivers version is 4.60.17718. The CIM provider ver- sion is 4.60.17718.
	Generic	 Added support for a new report for bad cables *** Fixed random parsing failures of string registry entries Fixed compilation failure of "Hel- lo_world" in the SDK Fixed the return value of ib_que- ry_ca() if failed to allocate resources for operation 	*** Requires firmware v2.30.8000 and above.
	Perfor- mance	 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 	

Release	Category	Description	Notes
	ND	 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 	

Table 20 - Change Log History

Release	Category	Description	Notes
	Ethernet	 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 	
	WMI/CIM	 Added ControlledBy association to IBPort Fixed ConformsToProfiles associa- tion for SoftwareIdentity and Drive- rIdentity Fixed execution of all tests which were running when executing Diag- nostic tests on one instance Fixed a failure occurred when run- ning MLNX_Card Fixed the printing of diagnostics log Fixed an issue preventing from get- event to show information after dis- abling the PCI device Removed support for the following configuration: ModeFlags SingleMsixNum MultiMsixNum SingleEqNum MultiEqNum DebugLevel DebugFlags UsePrio NumFcExch EnableQoS BlockMcastLoopBack InterruptFromFirstPacket ProbeVf 	

Release	Category	Description	Notes
4.55	Generic	 Added support for Windows Server 2012 R2 Operating System Added the ParentBusPath option to each port registry key Added a new hardware ID for Con- nectX®-3 Pro NICs The QP numbers allocation is now round-robin manner RecvCompletionMethod as Interrupt is no longer supported Removed the LsoV11Pv4 from the registry/UI Removed from the bus driver con- figuration the 'Non-DMA' option Removed the TXRingNum option from the UI 	

Table 20 - Change Log History

Release	Category	Description	Notes
	NVGRE	 Added NVGRE hardware off-load support (for ConnectX®-3 Pro cards only) Added to the UI the *Encapsulat- edPacketTaskOffload option when using ConnectX®-3 Pro NICs 	
	Perfor- mance	 Added the nd_send_bw and nd_send_lat ND benchmarking tools Fixed nd_*_bw to achieve better per- formance (memory buffer align- ment) and consistent results 	
	Ethernet	 Fixed the issue preventing messages to be sent in VLAN 0 when using many VMQ rings Added IP-IP checksum off-load sup- port Added Ports TX arbitration/Band- width allocation per port The following ND providers, MLX- 4ND and MLX4ND2 are installed by default Fixed setting the correct SL in UD traffic over RoCE 	
	InfiniBand	 IPoIB performance improvements Fixed a part_man issue related to wrong statistics over virtual partman interfaces 	
	RoCE	 Enabled roce_mode value overwrite in case it exists during installation Fixed in ibv_devinfo the display of correct transport RoCE mode Added Sniffer for RoCE packets The used RoCE mode set upon driver load is printed into event log message 	

Table 20 - C	Change Log	History
--------------	------------	---------

Release	Category	Description	Notes
4.40	Generic	 Added a notification in the event log in case SMB is not supported in ConnectX®-2 firmware Added the trace tool for WPP trac- ing Added copyright to the SDK files Added WMI/Powershell support Fixed an issue causing the setup to fail upon perf_tuning failure during the installation. An error message will be printed in the instal- lation log upon perf_tuning failure. Removed port setting registry key during uninstall Fixed and issue with the Mellanox adapter being shown on the USB removal menu, which caused the removal of the Mellanox adapter once removing the USB. 	
4.40	Perfor- mance	 Set 512 RX buffers by default Removed TXRingNum Changed the perf_tuning setting to achieve a better performance tuning Added the nd_write_bw/ nd_write_lat and nd_read_bw/ nd_read_lat tools Fixed the perf_tuning indication of the last chosen tuning scenarios Fixed a crash in the ib_send_lat/ bw utilities caused when the port link was down Fixed the "Restore to defaults" option in the perf_tuning tool. Now the default values are being restored 	

Release	Category	Description	Notes
	Ethernet	 Added Transmit Side Scaling (TSS) Added Ethernet QoS proprietary counters, diagnostics and traffic for monitoring, using Windows' perf- mon 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: Low Latency Moderate Aggressive In addition to old values that are not sup- ported 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 configura- tion on NICs with a single ports Fixed the values for the MTU and rate of the CM-REQ Fixed miniport reset on sending sce- narios 	
		disabling QoS	

Release	Category	Description	Notes
Release	Category Ethernet	 Description Enabled MaxRssProcessirs support of the following values: 1, 2, 4, 8, 16, 32, 64 Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation Network Direct: Fixed the swap of InboundReadLimit and Out- boundReadLimit when creating an EndPoint and in Connector::Get- ConnectionData Network Direct: Fixed disallowing creation of EndPoint with zero attri- butes 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 the MTU size the IP header size (1500 ->1514, 9600- >9614). Thus the minimum Jumbo frame size is 614. 	Notes
		 Interrupt moderation supports the following profiles: Low Latency Moderate Aggressive In addition to old values that are not supported anymore. Made mlx4_bus and Ethernet devices removable 	

Release	Category	Description	Notes
Kelease	Ethernet	 Description Network Direct: Added support for NDv2 Network Direct: Set the default ND provide value to mlx4nd2 Fixed WoL support on NIC with a single port Fixed the default RoCE configuration on NICs with a single ports Fixed the values for the MTU and rate of the CM-REQ Fixed miniport reset on sending scenarios Removed the QoS attributes when disabling QoS Enabled MaxRssProcessirs support of the following values: 1, 2, 4, 8, 16, 32, 64 Network Direct: Fixed a crash occurred when more than 4 SGEs elements were used in an ND write operation Network Direct: Fixed the swap of InboundReadLimit and OutboundReadLimit when creating an EndPoint and in Connector::GetConnectionData Network Direct: Fixed disallowing creation of EndPoint with zero attributes in the Receive Queue Network Direct: Removed the option of NDK registration failure requiring a reboot of the machine to register it again Network Direct: Added the option of sensing the incoming Read messages according to the device capabilities when creating an EndPoint limit 	•
		ND connectivity between VMs on the same host	

Release	Category	Description	Notes
	InfiniBand	 On rare occasions, depends on the GUID assignment, the IPoIB MAC address can be assigned with a multicast MAC (the least significant bit of the most significant address octet is set to 1). In that case, all of the traffic over the IPoIB I/F is dropped. If you experience this issue, please contact Mellanox support. Added active_mtu field to struct ib_port_attr_t Added the option of vstat displaying the active_mtu of the ports Allowed registration of a large Memory Region which is splitted to many segments Fixed a bluescreen issue that occurred when disabling the interface after a TX stress over the VMQ Fixed a failure of MPI/ND over InfiniBand Added the option of part_man printing the adapter name when the Port GUID is set to zero. Added the option of part_man printing the leading zeroes of port GUID Prevented displaying a message to 	
	installa- tion/ Upgrade	 Prevented displaying a message to upgrade the firmware for OEM NICs if it has the latest firmware version Removed portsetting registry key during uninstallation 	
4.3	Generic	• Added support for a new provider called MLX4ND, which supports both NDv1 and NDv2 interfaces	WinOF VPI version 4.3 was released as an intermediate release.

Table 20 -	Change I	Log History
------------	----------	-------------

Release	Category	Description	Notes
	Perfor- mance	 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: 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	 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 fol- lows: 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	• Added ibdiagnet utility support	

Release	Category	Description	Notes
4.2	Generic	 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 pro- cessors (more than 64 processors support) to allow assignment of MSI-X affinity for multiple proces- sor groups. Set an adequate number of MTTs to map all physical memory Allocated firmware and ICM mem- ory in chunks of non-paged memory instead of using contiguous physical memory. Fixed RSS indirection table map- ping building when there are less RX rings than RSS cores. Fixed a bug, preventing standard work with BAR value more than 4GB. Fixed error flows causing a Blue- screen in driver startup/unload Fixed a Bluescreen occurrence upon shutdown due to leak in active resources Changed device names in device manager and their hardware IDs. The changes were made to distin- guish between ConnectX®-2 and ConnectX®-3: for ConnectX-2: MLX4\ConnectX- 2_Eth and IBA\ConnectX-3_IPOIB for ConnectX-3: MLX4\ConnectX- 3_Eth and IBA\ConnectX-3_IPOIB Set QoS settings only for ConnectX- 3_Eth and IBA\ConnectX- 3_IDI and IBA\ConnectX- 3_IDIB 	

Table 20 - C	Change Log	History
--------------	------------	---------

Release	Category	Description	Notes
4.2	Generic	 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 TxRing-Num registry key from the UI Improved the "Port Protocol" dialog Added Registry key documentation to the setup package 	
	Performance	 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. Blue- Flame is now the default working mode for all packets that have a descriptor which fits into a BF regis- ter (currently 256 bytes). Use "Blue- Flame" 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 perfor- mance Added a performance UI to tune per- formance under various scenarios Added a tool to tune performance under various scenarios 	

Etherne	 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 	
---------	--	--

5 API Change Log History

Release	Name	Description
4.80.50000	RDMA_TRANSPORT_RDMAOE_1	It is an alias to: RDMA_TRANSPORT_RDMAOE
	RDMA_TRANSPORT_RDMAOE_1_25	Added enumerated values
	<pre>is_roce(), is_mac_based_roce(), is_ip_based_roce(), is_rro- ce_or_ip_based_roce()</pre>	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
	<pre>ib_get_mad_inner()</pre>	Changed API (one more input parameter was added)
	<pre>ib_get_mad()</pre>	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
4.40	VERBS_MINOR_VER	Increased its value, 0x0009 -> 0x000a
	enum eth_link_speeds	Added enumerated values
	struct ib_port_attr_t	 The mtu field was separated into two fields: 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