



# Mellanox WinOF VPI Archived Known Issues

---

Last Updated: 28 January 2018

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 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

© 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

Table of Contents .....	3
List Of Tables .....	4
Chapter 1 Archived Known Issues .....	5

## List Of Tables

Table 1: Archived Known Issues .....	5
--------------------------------------	---

# 1 Archived Known Issues

The following table describes archived known issues in the former releases on WinOF.

**Table 1 - Archived Known Issues (Sheet 1 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> Pinning all the physical memory (used by RDMA operations, such as register memory, pin user memory) on the machine, on Operating Systems prior to Windows Server 2012, may cause the machine to hang.</p> <p><b>Workaround:</b> Avoid pinning the whole machine memory in those Operating Systems.</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> 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.</p> <p><b>Workaround:</b> Close all applications that use ND or libibumad before performing shutdown/restart/hibernate.</p> <p><b>Keywords:</b> Generic</p>
322721	<p><b>Description:</b> Activating NC-SI in WinOF v4.90.10541 may cause driver's loading failure when using an older firmware version than 2.30.8000.</p> <p><b>Workaround:</b> Do not enable NC-SI in machines that WinOF v4.90.10541 is installed in.</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> The maximum values returned by the <code>ib_query_ca()</code> function (for example: <code>max_qp</code>, <code>max_mr</code>) are the upper limits of the supported resources by the device. However, it may be impossible to use these maximum values, since the actual number of any resource that can be created may be limited by the machine configuration, the amount of host memory, user permissions, and the amount of resources already in use by other users/processes.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> Running <code>Ntttcp</code> without the “-a X” flag (<math>X &gt; 1</math>) in a NIC configured with 10GbE, may cause low bandwidth in TCP single stream.</p> <p><b>Workaround:</b> Run <code>Ntttcp</code> with “-a 8” for best performance</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> Active links disappear after changing the cable connectivity from Ethernet to InfiniBand or vice versa.</p> <p><b>Workaround:</b> Disable and enable the <code>mlx4_bus</code> interface from the Device Manager.</p> <p><b>Keywords:</b> Generic</p>

**Table 1 - Archived Known Issues (Sheet 2 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> On ConnectX®-2/ConnectX®-3 Ethernet adapter cards, there is a mismatch between the GUID value returned by firmware management tools and that returned by fabric/driver utilities that read the GUID via device firmware (e.g., using ibstat). Mlxburn/flint return 0xffff as GUID while the utilities return a value derived from the MAC address. For all driver/firmware/software purposes, the latter value should be used.</p> <p><b>Workaround:</b> Use the GUID value returned by the fabric/driver utilities (not 0xffff).</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> The UI displays the network interface’s state incorrectly even after disabling and re-enabling the "disconnected" network interfaces.</p> <p><b>Workaround:</b> To see the correct state and link speed, perform one of the following:</p> <ul style="list-style-type: none"> <li>• Run Get-netadapter Powershell cmdlet</li> </ul> <p>or</p> <ul style="list-style-type: none"> <li>• Right click on that interface from “Network Connections” and click on status</li> </ul> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> On rare occasions, as a result of port configuration change (IB/ETH) the UI may get stuck for up to a few minutes. This effect does not require any user action. The UI returns to its proper functionality after a few minutes.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> The drivers’ functionality is limited up to 128 cores.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Generic</p>
-	<p><b>Description:</b> According to the hardware architecture, Mellanox NIC devices are presented in the registry as virtual devices and not as physical devices.</p> <p><b>Workaround:</b> Any customer programs or scripts that look into the NIC registry values should not assume it is a physical device.</p> <p><b>Keywords:</b> Generic</p>
698291	<p><b>Description:</b> A BSOD can happen when the driver is enabled and disabled in loops and without any pause.</p> <p><b>Workaround:</b> Disable the driver only after the Ethernet or IPoIB adapter is started. Starting any of these adapters can be checked by pinging them.</p> <p><b>Keywords:</b> Generic</p>

**Table 1 - Archived Known Issues (Sheet 3 of 18)**

Internal Ref.	Issue
464449	<b>Description:</b> IPv6 traffic between Hyper-V hosts over IPoIB v-Switch may experience traffic loss.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
491546	<b>Description:</b> 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.
	<b>Workaround:</b> Remove the virtual port before the upgrade and re-add it after the upgrade.
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> InfiniBand application that was compiled with an SDK version earlier than WinOF v4.90 is not binary compatible.
	<b>Workaround:</b> 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.
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> Without separate ports for each stream, WinSock multiplexes every packet to every subscriber socket and then filters it out.
	<b>Workaround:</b> Use different UDP ports to get higher performance when using multi-cast packets.
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> 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.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> The “Packets Received Discarded” and “Packets Received Errors” counter may display wrong results.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> Connection failure on ND tests while machine A have IBAL provider and machine B have MLX4ND provider.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand

**Table 1 - Archived Known Issues (Sheet 4 of 18)**

Internal Ref.	Issue
-	<b>Description:</b> Hibernate and Sleep are not functional when user-space is using its resources.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> IPoIB does not support: <ul style="list-style-type: none"> <li>• MAC address change</li> <li>• QoS (packet priority)</li> <li>• Connected Mode</li> </ul>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> 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.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> OpenSM does not run as a service during installation since the SM requires the GUID parameter to decide which port to work on. Setting it on setup causes it to work only on the first port and not the others.
	<b>Workaround:</b> To run OpenSM as a service, assuming the package was installed in the default path, use: <pre>sc create opensm binPath= "c:\Program Files\Mellanox\MLNX_VPIB\Tools\opensm.exe"</pre> To start the service run: <pre>sc start opensm</pre>
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> Tools issues: <ul style="list-style-type: none"> <li>• ibportstate does not work on RoCE ports</li> <li>• ibdiagpath may crash on Hyper-V machines</li> </ul>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> If an application which uses InfiniBand runs while the driver is being restarted, a bluescreen or an NMI may occur.
	<b>Workaround:</b> Stop all InfiniBand applications including OpenSM upon driver restart.
	<b>Keywords:</b> InfiniBand



**Table 1 - Archived Known Issues (Sheet 5 of 18)**

Internal Ref.	Issue
330284	<b>Description:</b> 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.
	<b>Workaround:</b> To resume operation, stop OpenSM and restart the driver while OpenSM is down.
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> No communication between the physical interface and a VM that uses vSwitch created over virtual IPoIB (Pkey), and vice versa.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
90224	<b>Description:</b> Burning a firmware image with a "mtu_cap" value other than the default causes the driver load to fail.
	<b>Workaround:</b> Do not change the "mtu_cap" value
	<b>Keywords:</b> InfiniBand
383460	<b>Description:</b> Calling <code>ib_query_cq()</code> on a CQ which was created with 1 entry indicates that there are 0 entries in the CQ.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand
439805	<b>Description:</b> IPoIB ports report the configured MTU instead of the effective MTU. These MTUs differ when the corresponding InfiniBand partition is configured with a smaller MTU than the port’s configured MTU. In such case, there may be communication failures and/or CPU work overhead on the communicating nodes (even if the peer node is configured correctly).
	<b>Workaround:</b> Make sure the configured adapter MTU (the “Jumbo packet” property in the adapter’s advanced settings) is consistent with the corresponding InfiniBand partition MTU. Specifically, the default IPoIB adapter MTU is 4KB and the default OpenSM partition MTU is 2KB. Thus, the default MTU of adapters and OpenSMs is inconsistent. Therefore, the MTU of either all adapters or all OpenSMs (which may be on network switches) on the InfiniBand subnet needs to be changed. Note that the adapter's MTU can be configured through PowerShell.
	<b>Keywords:</b> InfiniBand
-	<b>Description:</b> The displayed MAC address in the DHCP server of virtual IPoIB I/F may display a wrong data (FF00.....) although the I/F is still fully functional.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> InfiniBand

**Table 1 - Archived Known Issues (Sheet 6 of 18)**

Internal Ref.	Issue
498130	<p><b>Description:</b> 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.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> InfiniBand</p>
-	<p><b>Description:</b> Disabling the “Priority &amp; VLAN tag” in the UI which VLANID is configured, may result in sending packets with the configured VLANID.</p> <p><b>Workaround:</b> Remove the VLANID before disabling the “Priority &amp; VLAN tag”.</p> <p><b>Keywords:</b> Ethernet</p>
206528/ 206945	<p><b>Description:</b> WakeOnMagicPacket registry key is not added to the registry although WoL is supported by the driver and by the NIC.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet</p>
189704	<p><b>Description:</b> When the ports of the device are configured as Ethernet only, ibstat/vstat may display wrong information.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet</p>
-	<p><b>Description:</b> High multicast drop rate on multicast storming.</p> <p><b>Workaround:</b> Use “Multicast traffic” tuning option under the performance tab. For further information, please refer to section “Tunable Performance Parameters” in the User Manual.</p> <p><b>Keywords:</b> Ethernet</p>
-	<p><b>Description:</b> When there is a stress in TCP connection establishments, some of those connections may fail.</p> <p><b>Workaround:</b> Increase the Ring queue sizes:</p> <ul style="list-style-type: none"> <li>• ReceiveBuffers - controls the receive ring size</li> <li>• TransmitBuffers - controls the transmit ring size</li> </ul> <p><b>Keywords:</b> Ethernet</p>

**Table 1 - Archived Known Issues (Sheet 7 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> The DCB component specifies a default traffic classification that is applied to all egress packets that do not match other classification conditions. In this case, the network adapter assigns the IEEE 802.1p priority level that is associated with the default classification to these egress packets. The default traffic classification has the following attributes:</p> <ul style="list-style-type: none"> <li>• It has a traffic classification condition of type NDIS_QOS_CONDITION_DEFAULT.</li> <li>• It is the first traffic classification defined in the array of NDIS_QOS_CLASSIFICATION_ELEMENT structures.</li> </ul> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet</p>
410377	<p><b>Description:</b> Virtual Ethernet Interfaces created by vea_man are not tuned by the automatic performance tuning script.</p> <p><b>Workaround:</b> For optimal performance need to follow the performance tuning guide and apply relevant changes to the VEA interface.</p> <p><b>Keywords:</b> Ethernet</p>
326885	<p><b>Description:</b> Wake-on-Lan (WoL) cannot be disabled on NICs which supports it.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet</p>
687328	<p><b>Description:</b> When SR-IOV is turned on and the “Jumbo Packet” field under the “Advanced” tab of the Ethernet driver's properties (in Device Manager) is set to 9602 or more, IP fragmentation does not work.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Ethernet</p>
-	<p><b>Description:</b> Running Quality of Service (QoS) commands without the parameter “-PolicyStore ActiveStore” may cause machines to load without Quality of Service policy.</p> <p><b>Workaround:</b> Store the QoS policy in the ActiveStore</p> <p><b>Keywords:</b> Quality of Service</p>
-	<p><b>Description:</b> RoCE does not support:</p> <ul style="list-style-type: none"> <li>• Traffic cannot go through the router. It works in the same subnet only</li> <li>• Multicast traffic</li> <li>• VLAN</li> <li>• Layer 3 feature</li> </ul> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>

**Table 1 - Archived Known Issues (Sheet 8 of 18)**

Internal Ref.	Issue
327531	<p><b>Description:</b> In machines with heterogeneous NICs, a NIC which supports RoCE v2, and a NIC which does not support RoCE v2 the following issues might raise:</p> <ul style="list-style-type: none"> <li>• ConnectX®-3 Pro is loaded with the transport type RoCE v2</li> <li>• ConnectX®-3 Pro is loaded with the transport type RoCE</li> </ul> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>
-	<p><b>Description:</b> When NVGRE off-load is enabled, the GRE traffic cannot be accepted as a regular L2 traffic and requires special L2_TUNNELING steering rules. In such case the GRE packets are dropped or directed to promiscuous queue.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>
193619	<p><b>Description:</b> Using different versions of RoCE in your cluster is not supported.</p> <p><b>Workaround:</b> Use the same RoCE version in all the cluster in the Ethernet ports.</p> <p><b>Keywords:</b> RoCE NVGRE</p>
-	<p><b>Description:</b> RDMA Activity counters do not count during NetworkDirect RoCE traffic.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>
-	<p><b>Description:</b> GRE traffic steering by inner MAC and by outer MAC simultaneously is currently not supported.</p> <p><b>Workaround:</b> Configure steering or by inner MAC, or by outer MAC.</p> <p><b>Keywords:</b> RoCE NVGRE</p>
-	<p><b>Description:</b> 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.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>
-	<p><b>Description:</b> Set the bus driver registry keyAcceptGREbyOuterMAC_P1/2 per port to accept GRE traffic by outer MAC and to duplicate L2 steering rule to L2_TUNNELING rule for each VMQ set filter request without GRE flag.</p> <p><b>Note:</b> For regular NVGRE Hyper-V scenarios the value of the registry key below must be set to 0:AcceptGREbyOuterMAC_P1/2</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE NVGRE</p>

**Table 1 - Archived Known Issues (Sheet 9 of 18)**

Internal Ref.	Issue
206696	<p><b>Description:</b> When using WinOF 4.40 or above, low throughput will be seen on 40GbE adapters when QoS is enabled.</p> <p><b>Workaround:</b> Disable QoS when it is not in use. Open a PowerShell prompt. Run: Disable-NetAdapterQos -name &lt;Interface Name&gt; where &lt;Interface Name&gt; is e.g. "Ethernet 1"</p> <p><b>Keywords:</b> Performance</p>
-	<p><b>Description:</b> perf_tuning is supported only when one of the two NUMA nodes are in use.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Performance</p>
-	<p><b>Description:</b> Running performance benchmarks for a short period of time (&lt; 1 sec) may provide bad latency in IPoIB and Ethernet.</p> <p><b>Workaround:</b> Set "Rx Interrupt Moderation Profile" and "Tx Interrupt Moderation Profile", to "Low Latency" to avoid bad latency.</p> <p><b>Note:</b> This may increase CPU utilization.</p> <p><b>Keywords:</b> Performance</p>
-	<p><b>Description:</b> The driver uses optimal interrupt moderation values for 10 GbE SR-IOV VF scenario. For other scenarios, the optimal values yet to be found.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Performance</p>
443137/ 439897	<p><b>Description:</b> 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.</p> <p><b>Workaround:</b> Reduce the counters sample rate.</p> <p><b>Keywords:</b> Performance</p>
-	<p><b>Description:</b> When the vSwitch is detached from the ETH\IPoIB device while the driver is disabled, the device does not reacquire the static IP it had before the attachment of the vSwitch. When the vSwitch is attached to the ETH\IPoIB device while there is no link, it will not receive the device IP when the link is back up.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>

**Table 1 - Archived Known Issues (Sheet 10 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> After attaching the vSwitch to the ETH\IPoIB device, changing the “Jumbo Packet” registry key on the ETH\IPoIB device does not affect the vSwitch configuration and vice versa. For example, if the user sets the “Jumbo Packet” on the ETH\IPoIB device to X, and the “Jumbo Packet” on the vSwitch to X+Y, X+Y sized packets will be passed from NDIS down to the driver and they will be dropped by it.</p> <p><b>Workaround:</b> Reattach the vSwitch to sync with the value set in the ETH\IPoIB device.</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> Unexpected behavior might occur when running in a virtualized environment and creating two virtual switches bound to each of the ports of a dual port NIC and then using both of them with two vNICs from the same VM.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> In IPoIB when using long Multicast traffic from a Virtual Machine (VM) to an external host there might be up to 0.5% loss in 5% bursts.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> Hyper-V is at low bandwidth on LBFO vSwitch, Windows Server 2012.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> In Ethernet to achieve better iperf TCP performance between a Linux VM and a Windows VM on different hosts, when using MS MUX over the Ethernet driver, use the non VMQ mode for the VMs.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> After disabling and enabling a port on a guest, a ping to it may be renewed after a minute. The ARP requests sent by Windows are less frequent as the time passes. If the guest port was down for a while, it could take time until Windows decides to send another ARP request to it.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>

**Table 1 - Archived Known Issues (Sheet 11 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> When VMQ is enabled after reset, the driver loads all the VMQs that existed before the reset. However, it is not guaranteed that each VMQ will receive the same QP number it had before the reset. This can cause some delay as a result of resetting before connectivity is reestablished. The delay is caused by the time it takes for the ARP table to update after initiating the Gratuitous ARP.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> The IPoIB non-VMQ mode is supported only when the VMQ is enabled according to the registry values.</p> <p><b>Workaround:</b> To use the non-VMQ mode for a VM, change its settings as follow:</p> <ul style="list-style-type: none"> <li>• Press “Settings” on the VM</li> <li>• Go to Network Adapter -&gt; Hardware Acceleration</li> <li>• Un-check the “Enable virtual machine queue”</li> </ul> <p><b>Keywords:</b> Hyper-V</p>
-	<p><b>Description:</b> Working with Windows guest OS over non-Windows SR-IOV hypervisor may result with higher latency compared to Windows Hypervisor.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> SR-IOV</p>
-	<p><b>Description:</b> When working with OEM adapter cards in SR-IOV mode, the VF is identified as a Mellanox adapter. There is no other effect on the behavior of the VF.</p> <p><b>Workaround:</b> Verify that the FW INI contains the parameter <code>vf_subsystem_id</code> with the OEM encoding in the [HCA] section.</p> <p><b>Keywords:</b> SR-IOV</p>
-	<p><b>Description:</b> An SR-IOV Virtual Machine is loaded in InfiniBand mode when no OpenSM is available in the subnet. The following event appears in the event log: “detected a null port GUID for port &lt;NUMBER&gt;. A Virtual Function device may have a null port GUID if there is no OpenSM instance on its network. Please make sure the network has an active OpenSM and restart the driver.”</p> <p><b>Workaround:</b> Start the OpenSM and restart the driver.</p> <p><b>Keywords:</b> SR-IOV</p>
418268	<p><b>Description:</b> For InfiniBand SR-IOV guest, OpenSM Assigned GUIDs are not supported and may cause unexpected behavior.</p> <p><b>Workaround:</b> Work only with Administrator assigned GUIDs.</p> <p><b>Keywords:</b> SR-IOV</p>

**Table 1 - Archived Known Issues (Sheet 12 of 18)**

Internal Ref.	Issue
408734	<b>Description:</b> 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.
	<b>Workaround:</b> Disable the Para-Virtualization adapter.
	<b>Keywords:</b> SR-IOV
-	<b>Description:</b> In SR-IOV mode over Hyper-V, all ports are in pure Ethernet mode and RDMA is not supported on either port.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV
342421	<b>Description:</b> 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.
	<b>Workaround:</b> Unbond all Vswitches from the NIC's interfaces before enabling SR-IOV in the BIOS.
	<b>Keywords:</b> SR-IOV
427038	<b>Description:</b> In SR-IOV, 40Gbps interfaces of vSwitch interface on HyperV, and the VMNIC are identified as 10Gbps. Despite the incorrect speed identification, they still achieve 40Gbps performance.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV
-	<b>Description:</b> Device name in the Device Manager does not show the correct OEM branding for SR-IOV Virtual Function devices without the updated firmware.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV
759813	<b>Description:</b> When working in SR-IOV mode, if the version of the driver installed on the hypervisor is v5.02 or v5.22, the driver installed on the VM must also be v5.02 or 5.22 and above.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> SR-IOV
-	<b>Description:</b> Upgrading the driver while the UI is opened with the “ConnectX NIC device” may cause the installation process to never end.
	<b>Workaround:</b> Close the UI before driver upgrade.
	<b>Keywords:</b> Installation / Upgrade



**Table 1 - Archived Known Issues (Sheet 13 of 18)**

Internal Ref.	Issue
207497	<b>Description:</b> Rebooting the machine while uninstalling WinOF may result in installation failure.
	<b>Workaround:</b> Delete Mellanox components from HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\DIFxApp\Components. The Mellanox components are mlx4eth63, ipoib6x and mlx4_bus
	<b>Keywords:</b> Installation / Upgrade
137859	<b>Description:</b> Canceling the installation process may leave the bus driver in a disable state. The driver appears in a yellow bang containing the following error message: "Windows cannot start this hardware device because its configuration information (in the registry) is incomplete or damaged. (Code 19)".
	<b>Workaround:</b> Uninstall the current version and install the older one.
	<b>Keywords:</b> Installation / Upgrade
198537	<b>Description:</b> Downgrade is not supported.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> Configuration can be restored only in Windows Server 2012 and above.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> Running a downgrade in silent mode is not supported. Upon downgrade the return code will always be 0.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> Uninstalling the driver when multiple of VLANs are configured never ends.
	<b>Workaround:</b> Remove the VLANs before uninstallation.
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> Upgrading the driver while Perfmon is open and monitoring the "ConnectX NIC device" may end up with adapters having yellow bang with code 10.
	<b>Workaround:</b> Close Perfmon before driver upgrade.
	<b>Keywords:</b> Installation / Upgrade

**Table 1 - Archived Known Issues (Sheet 14 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> Installation/upgrade fails due to PNP failure to copy the driver files to the driver store, and the following text is printed in the event logs:            Fault bucket, type 0            Event Name: PnPDriverImportError            Response: Not available            Attached files:            C:\Users\<user>\AppData\Local\Temp\DMI151A.tmp.log.xml            C:\Program Files\Mellanox\MLNX_VPI\ETH\mlx4eth63.inf</user></p> <p><b>Workaround:</b> Reboot the machine and reinstall</p> <p><b>Keywords:</b> Installation / Upgrade</p>
403352	<p><b>Description:</b> 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”</p> <p><b>Workaround:</b> Kill the WMIPrvSE.exe tasks in the task manager and reinstall.</p> <p><b>Keywords:</b> Installation / Upgrade</p>
-	<p><b>Description:</b> Following the upgrade of Mellanox driver to WinOF-4.60 and above or on servers with no Internet access, the first PowerShell command might be stuck for ~2-3 minutes before its completed.            According to the following content, this issue is related to .Net framework version or an issue with the Internet access:  <a href="http://www.minasi.com/forum/topic.asp?TOPIC_ID=39253">http://www.minasi.com/forum/topic.asp?TOPIC_ID=39253</a></p> <p><b>Workaround:</b> Run the following script on the server to optimizes loading PowerShell DLLs:  <pre> \$Env:PATH = [Runtime.InteropServices.RuntimeEnvironment]::GetRuntimeDirectory() [AppDomain]::CurrentDomain.GetAssemblies()   % { \$pt = \$_.Location if (! \$pt) {continue} if (\$cn++) {"} \$na = Split-Path -Leaf \$pt Write-Host -ForegroundColor Yellow "NGENing \$na" ngen install \$pt } </pre></p> <p><b>Keywords:</b> Installation / Upgrade</p>
417380/ 415257	<p><b>Description:</b> On ConnectX®-3 cards only, when upgrading from Windows Server 2012 R2 Inbox driver to WinOF, the RoCE mode setting in the registry is not properly transferred to the new driver. In case a non-default value was used it will not be configured following the upgrade.</p> <p><b>Workaround:</b> Reconfigure the RoCE Mode setting manually.</p> <p><b>Keywords:</b> Installation / Upgrade</p>

**Table 1 - Archived Known Issues (Sheet 15 of 18)**

Internal Ref.	Issue
427676	<b>Description:</b> Uninstall from the Device Manager is currently unsupported
	<b>Workaround:</b> Driver uninstall can be completed from the Programs and Features window as explained in the “Uninstalling Mellanox WinOF Driver” section in the User Manual.
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> WinOF Inbox driver does not support upgrade. When installing WinOF v4.40 and above on a Windows Server 2012 and above machine, the Inbox driver is uninstalled prior to starting the new installation and any previous configurations is lost. The Inbox driver will be reinstalled automatically when the new driver is uninstalled.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> Firmware upgrade may fail during installation if there was a prior firmware upgrade on the machine without a reboot after it. A firmware upgrade failure does not fail the whole installation.
	<b>Workaround:</b> Upgrade the firmware manually.
	<b>Keywords:</b> Installation / Upgrade
492398	<b>Description:</b> If there are disabled network interfaces or a disabled Mellanox bus driver, they will be enabled after the WinOF upgrade.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Installation / Upgrade
-	<b>Description:</b> ibdump may encounter packet drops upon a burst of more than 4096 (or $2^{\text{max-burst}}$ ) packets.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Utilities
-	<b>Description:</b> Packets loss is not reported by ibdump.
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Utilities
-	<b>Description:</b> Running ibdump on a RoCE//RoCE v1.5 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.
	<b>Workaround:</b> Verify Ethernet flow control is enable to ensure a lossless link
	<b>Keywords:</b> Utilities

**Table 1 - Archived Known Issues (Sheet 16 of 18)**

Internal Ref.	Issue
-	<p><b>Description:</b> Pcap file issues for RoCE IP Based:</p> <ul style="list-style-type: none"> <li>The packets 'capture-time' field is wrong (may be 0 or negative numbers).</li> <li>For every captured packet, an additional 0 size flame is added. This appears in Wireshark as a 'malformed Ethernet packet'.</li> </ul>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Utilities
-	<p><b>Description:</b> Sniffing over IB ports is currently not supported</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Utilities
199079	<p><b>Description:</b> When the tracer tool traces many events, it may consume a large amount of memory (up to several GB RAM).</p>
	<b>Workaround:</b> Use the button to set maximum number of the displayed lines at a maximum buffer size.
	<b>Keywords:</b> Utilities
367772	<p><b>Description:</b> ibdump works only on Ethernet links.</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> Utilities
-	<p><b>Description:</b> Running Microsoft CIM cmdlets operations and their derived classes on classes MLNX_NetAdapterSettingData and MLNX_NetAdapter-RoceSettingData is not supported. Calling those commands may cause the debugger, if connected to the machine, to assert.</p>
	<b>Workaround:</b> Use DriverCoreSettings instead.
	<b>Keywords:</b> CIM/WMI
408230	<p><b>Description:</b> For PCI Gen3, PcieLinkSpeed is reported as “Unknown” when running Get-NetAdapterHardwareInfo Powershell cmdlet</p>
	<b>Workaround:</b> N/A
	<b>Keywords:</b> CIM/WMI
-	<p><b>Description:</b> WMI does not work due to lack of permissions.</p>
	<b>Workaround:</b> Change the execution policy. Run: Set-ExecutionPolicy AllSigned
	<b>Keywords:</b> CIM/WMI
433986	<p><b>Description:</b> The information that is printed in the cmdlets get-netadaptersriov and Get-MlnxPCIDeviceSetting is inconsistent.</p>
	<b>Workaround:</b> Use only the cmdlet Get-MlnxPCIDeviceSetting.
	<b>Keywords:</b> CIM/WMI

**Table 1 - Archived Known Issues (Sheet 17 of 18)**

Internal Ref.	Issue
432674	<p><b>Description:</b> There is an interoperability problem between NDK and ND in RDMA operations during the write and read operations. However, the send operation resumes working. This happens since ND uses remote token in a network order (big endian) while NDK uses remote token in CPU order (little endian). Therefore, an inconsistency is caused between ND and NDK in RDMA operations.</p> <p><b>Workaround:</b> 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.</p> <p><b>Keywords:</b> ND</p>
-	<p><b>Description:</b> 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.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ND</p>
-	<p><b>Description:</b> Changing the default ND providers may cause random errors, such as: 0xC0000238 (NT_STATUS_ADDRESS_ALREADY_ASSOCIATED) on Connect() or with 0xC0000236 (NT_STATUS_CONNECTION_REFUSED) on Accept(). These errors can be safely ignored.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> ND</p>
-	<p><b>Description:</b> The driver fails to load on guest machine with Windows 10 and Inbox driver 4.91 when in SR-IOV mode.</p> <p><b>Workaround:</b> Since SR-IOV mode in v5.25 driver does not support Windows 10 old Inbox driver (4.91), we recommend you update your driver with the latest version of Windows 10.</p> <p>Keywords: SR-IOV, Windows 10</p>
-	<p><b>Description:</b> LBFO over IPoIB is not supported.</p> <p><b>Workaround:</b> N/A</p> <p>Keywords: LBFO, IPoIB</p>
648334	<p><b>Description:</b> The IPoIB driver may stall upon disabling, due to a non-zero reference count in TCP/IP stack.</p> <p><b>Workaround:</b> N/A</p> <p>Keywords: IPoIB, disable</p>

**Table 1 - Archived Known Issues (Sheet 18 of 18)**

Internal Ref.	Issue
558664	<p><b>Description:</b> In case a NIC with two ports (an Ethernet port and an InfiniBand port) is set to No Roce via config roce_mode, despite the configuration, RoCE will not be disabled for the Ethernet port.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RoCE</p>
664882	<p><b>Description:</b> Endure with MTT optimization for continuous memory allocation is not supported.</p> <p><b>Workaround:</b> N/A</p> <p><b>Keywords:</b> RDMA</p>