

Known Issues

For a list of old Know Issues, please see the relevant Release Notes version.

Internal Ref.	Issue
2724780	Description: On very rare cases a DevX call to create a native MKEY will fail due to fragmented memory in the allocated UMEM causing the UMEM page offset and the mkey page offset to misalign.
	Workaround: N/A
	Keywords: DevX, MKEY
	Detected in version: 2.70.51000
2727039	Description: WinOF-2 installation package will not automatically update the firmware on devices that are using secured firmware.
	Workaround: Manually install the firmware (see section " Firmware Upgrade " on how to).
	Keywords: Firmware upgrade, secure firmware
	Detected in version: 2.70.51000
2710916	Description: Wrong values on the VF-counters are exposed on the Hypervisor. "Packets Received Discarded" and "Packets Received Errors" of the counter-set "Mellanox WinOF-2 VF Port Traffic" represent values taken from the global-device or the PF specific.
	Workaround: N/A
	Keywords: Counters
	Detected in version: 2.70.50000
2683075	Description: MPReset handler may be triggered by the OS when using Windows Server 2022 due to some OIDs (e.g. OID_NIC_SWITCH_DELETE_VPORT) that can take a very long time to be completed.
	Workaround: Increase the CheckForHang timeout using our registry key "CheckForHangTOInSeconds" .
	Keywords: MiniportReset
	Detected in version: 2.70.50000
2701735	Description: Disabling one of the GPUs while the application is running could lead to system crash.
	Workaround: N/A
	Keywords: GPU
	Detected in version: 2.70.50000
2491846	Description: As oversubscription of QP parameters (entries and depth) is allowed, it could cause run-time failure when running out of resources.
	Workaround: N/A
	Keywords: QP creation
	Detected in version: 2.70.50000
2690140	Description: Requests of QPs with a string of values set to "max" (e.g., Max Queue Depth + Max SGE counter + Max inline Data size) cannot be processed by the driver as their accumulative size overcomes the WQ maximum size.
	Workaround: N/A
	Keywords: ND QP creation
	Detected in version: 2.70.50000
2380684	Description: Although the IPOIB failover team gets the correct DHCP address when first created, if the team is disabled and then enabled, Windows requests and rejects the DHCP address as BAD_ADDRESS.
	Workaround: When the issue is seen, restart the secondary member(s) of the team.
	Keywords: IPOIB teaming, DHCP
	Detected in version: 2.70.50000

2603423	Description: When in ETH mode, setting the MTU (JumboPacket) lower than 1514, results in Received Packets Error counters not being increased when receiving packets with larger frame size but less or equal to 1518 bytes (Like ping with data size of 1476).
	Workaround: N/A
	Keywords: MTU, traffic, counters
	Detected in version: 2.70.50000
2403963	Description: The DHCP redirect feature is not supported over FreeBSD VMs. When activated, DHCP packets will be dropped and VM will lose connectivity due to missing IP.
	Workaround: N/A
	Keywords: DHCP Redirect
	Detected in version: 2.60.50000
2397036	Description: On BlueField-2 setup, the maximum number of VFs enabled is less than the actual value supported by the firmware. When in SmartNIC mode, the number of VFs will decrease the SmartNIC enablements. When in separate mode, the number of supported VFs will be half of the firmware value as the VFs are split between the host and the Arm.
	Workaround: N/A
	Keywords: BlueField, VFs
	Detected in version: 2.60.50000
2374101	Description: After upgrade, *PtpHardwareTimestamp remains enabled. When *PtpHardwareTimestamp is enabled, UDP performance feature (URO) will be automatically disabled. This is an OS limitation, if you do not use the HW time stamp feature, it is recommended to disable this feature by setting *PtpHardwareTimestamp to 0.
	Workaround: Disable HW timestamping. by setting *PtpHardwareTimestamp to 0.
	Keywords: *PtpHardwareTimestamp, UDP performance feature ,URO
	Detected in version: 2.60.50000
2306807	Description: When the Decouple VmSwitch protocol is enabled, VM's friendly given name is not displayed when running the "Get-NetAdapterSriovVf" and "mlnx5hpccmd -DriverVersion" commands.
	Workaround: N/A
	Keywords: HPC, SR-IOV
	Detected in version: 2.60.50000
2205722	Description: WinOF-2 driver does not support IB MTU lower than 614.
	Workaround: N/A
	Keywords: IB MTU
	Detected in version: 2.60.50000
2180714	Description: In case the user configs TCP to priority 0 with no VlanID, the packets will be sent without a VLAN header since the miniport cannot distinguish between priority 0 with VlanId 0 and no Vlan tag.
	Workaround: N/A
	Keywords: TCP QOS
	Detected in version: 2.50.50000
2216232	Description: As ConnectX-5 adapter cards do not create counters for RX PACKET MARKED PCIe BUFFERS, its value will be 0.
	Workaround: N/A
	Keywords: ECN Marking
	Detected in version: 2.50.50000
2243909	Description: The driver to sends a wrong CNP priority counter while running RDMA.
	Workaround: Change the CNP priority using mlxconfig.

	Keywords: RDMA, CNP
	Detected in version: 2.50.50000
2118837	Description: Performance degradation might be experienced during UDP traffic when using a container networking and the UDP message size is larger than the MTU size .
	Workaround: N/A
	Keywords: Nested Virtualization, container networking
	Detected in version: 2.50.50000
2137585	Description: While working in IPoIB mode and *JumboPacket is set in the range of [256, 614], the driver issues a warning event log message (Event ID: 25). This is a false alarm and could be ignored.
	Workaround: N/A
	Keywords: JumboPacket
	Detected in version: 2.50.50000
2148077	Description: Explicitly disabling the *NetworkDirect key when using the HyperV mode, disables NDSPi as well as the NDK.
	Workaround: Enable NetworkDirect (ND).
	Keywords: ND, HyperV
	Detected in version: 2.50.50000
2117964	Description: A delay in connection establishment might be experienced when the ND application is started immediately after restarting the adapter card. This scenario occurs because the ND application requires the ARP table to find the destination MAC and generate the ARP request.
	Workaround: Use static ARP. Ping the system before starting the ND application.
	Keywords: ND, RDMA
	Detected in version: 2.40.51000
2117636	Description: On a native setup, when setting JumboPacket to be less than 1514, the Large Receive Offload (LRO) feature might be disabled, and all its counters will not be valid.
	Workaround: N/A
	Keywords: LRO, RSC
	Detected in version: 2.40.51000
2083686	Description: As PCIe Write Relaxed Ordering is enabled by default, some older Intel processors might observe up to 5% packet loss in high packet rate and small packets. (https://lore.kernel.org/patchwork/patch/820922/)
	Workaround: Disable the Relaxed Ordering Write option by setting the RelaxedOrderingWrite registry key to 0 and restart the adapter.
	Keywords: PCIe Write Relaxed Ordering
	Detected in version: 2.40.50000
1763379	Description: On Windows Server 19H1, running "netstat -axn" when RDMA is enabled and a vNIC is present, results in RDMA being disabled on the port with the VMswitch.
	Workaround: N/A
	Keywords: VMswitch, RDMA, Windows Server 2019
	Detected in version: 2.40.50000
1908862	Description: When running RoCE traffic with a different RoceFrameSize configuration, and the fabric (jumbo packet size) is large enough, the MTU will be taken from the initiator even when it supports larger size than the server.
	Workaround: N/A
	Keywords: RoCE, MTU
	Detected in version: 2.40.50000
1846356	Description: The driver ignores the value set by the "*NumVfs" key. The maximal number of VFs is the maximal number of VFs supported by the hardware.

	Workaround: N/A
	Keywords: SR-IOV NUMVFs
	Detected in version: 2.30.50000
1598716	Description: Issues with the OS' "SR-IOV PF/VF Backchannel Communication" mechanism in Windows Server 2019 Hyper-V, effect VF-Counters functionality as well.
	Workaround: N/A
	Keywords: Mellanox WinOF-2 VF Port Traffic, VF-Counters
	Detected in version: 2.30.50000
1702662	Description: On Windows Server 2019, the physical media type of the IPoB NIC will be 802.3 and not InfiniBand.
	Workaround: Use the mlx5cmd tool ("mlx5cmd -stat") which is part of the driver package to display the lin_layer type.
	Keywords: Windows Server 2019, IPoB NdisPhysicalMedium
	Detected in version: 2.20
1718201	Description: Heavy traffic causes Sniffer' limit file to be the same as the buffer size (100M by default).
	Workaround: N/A
	Keywords: Sniffer, heavy traffic
	Detected in version: 2.20
1576283	Description: When working with SR-IOV in Windows Server 2019, the vNIC that is working in SR-IOV mode status will be displayed as "Degraded (SR-IOV not operational)" although the SR-IOV VF is fully operational. The message can be safely ignored.
	Workaround: N/A
	Keywords: SR-IOV IB, Windows Server 2019
	Detected in version: 2.10
1580985	Description: iSCSI boot over IPoB is currently not supported.
	Workaround: N/A
	Keywords: iSCSI Boot, IPoB
	Detected in version: 2.10
1536971	Description: The RscIPv4 and RscIPv6 keys' values are set to 0 for the host in Windows Server 2019. As the values for those keys are already written by the Inbox Driver in Windows Server 2019, they will not be changed when upgrading.
	Workaround: N/A
	Keywords: RscIPv4, RscIPv6, Windows Server 2019
	Detected in version: 2.10
1419597	Description: On servers with large number of VMs, (typically more than 40), after restarting the NIC on the host, VMs' IPv6 global address is not retrieved back from the DHCP.
	Workaround: Restart the NIC inside the VM.
	Keywords: VMQ, SR-IOV
	Detected in version: 2.10
1419597	Description: On servers with a large number of VMs (typically > 40) - after a NIC restart on the host, VMs' IPv6 global address cannot be retrieved from DHCP.
	Workaround: Restart Microsoft NIC inside the VM.
	Keywords: VM, IPv6 address, DHCP
	Detected in version: 2.0
1336097	Description: Due to an OID timeout, the miniport reset is executed.

	<p>Workaround: Increase the timeout value in such way that $2 * \text{CheckForHangTOInSeconds} > \text{Max OID time}$.</p> <p>For further information, refer to section General Registry Keys in the User Manual.</p> <p>Keywords: Resiliency</p> <p>Detected in version: 1.90</p>
1310086	<p>Description: Multicast packets are passed via to the VM the Hyper-V (even in SR-IOV VMs). As such, the Hyper-V can decide to drop the packets based on its specific policy.</p> <p>Note: This issue is only related to FreeBSD OSes.</p> <p>Workaround: N/A</p> <p>Keywords: Hyper-V OS</p> <p>Detected in version: 1.90</p>
1154447	<p>Description: Adding diagnostic counters to performance monitor might cause counters to get cleared every several seconds.</p> <p>Workaround: Change the time period between samples to more than 1 second.</p> <p>Keywords: Diagnostic Counters</p> <p>Detected in version: 1.90</p>
1074589	<p>Description: When PXE boot is using Flexboot, the IPoIB interface is not receiving the reserved address from the DHCP using GUID reservation.</p> <p>Workaround: To obtain the reserved address, use a 6-byte MAC address instead of the 8-byte client ID.</p> <p>Keywords: PXE boot, IPoIB, Flexboot, DHCP</p> <p>Detected in version: 1.80</p>
917747	<p>Description: VF driver initialization fails in case of bad MSIX mapping when running Windows Server 2012 R2 Hypervisor with Windows Server 2016 VM with more than a single core CPU. As a result, performance desegregation might occur.</p> <p>Workaround: Run either with one CPU core, or run with different Operating Systems.</p> <p>Keywords: SR-IOV</p> <p>Detected in version: 1.80</p>
1170780	<p>Description: The driver must be restarted in order to switch from RSS to NonRSS mode. Therefore, if a PowerShell command is used on a specific VM to an enabled/disabled VMMQ without restarting the driver, the RSS counters will keep increasing in Perfmon.</p> <p>Workaround: Restart the driver to switch to NonRSS mode.</p> <p>Keywords: RSS, NonRSS, VMMQ</p> <p>Detected in version: 1.80</p>
1149961	<p>Description: In RoCE, the maximum MTU of WinOF-2 (4k) is greater than the maximum MTU of WinOF (2k). As a result, when working with MTU greater than 2k, WinOF and WinOF-2 cannot operate together.</p> <p>Workaround: N/A</p> <p>Keywords: RoCE, MTU</p> <p>Detected in version: 1.80</p>
1145421	<p>Description: In IPoIB SR-IOV setup, in the Hyper-V Manager, the address appears as "SR-IOV enabled" instead of "SR-IOV active". This does not influence any activity or functionality.</p> <p>Workaround: N/A</p> <p>Keywords: IPoIB SR-IOV setup, Hyper-V</p> <p>Detected in version: 1.80</p>
1145421	<p>Description: In the "Network Connections" panel of Virtual Function (VF) in IPoIB SR-IOV setup, the Microsoft adapter may appear in addition to the NVIDIA® adapter. This does not influence any activity or functionality.</p> <p>Workaround: N/A</p> <p>Keywords: Network Connections, VF, IPoIB SR-IOV</p>

Detected in version: 1.80

SR-IOV Support Limitations

The below table summarizes the SR-IOV working limitations, and the driver's expected behavior in unsupported configurations.

WinOF-2 Version	NVIDIA® ConnectX®-4 Firmware Ver.	Adapter Mode		
		InfiniBand	Ethernet	
		SR-IOV On	SR-IOV Off	SR-IOV On /Off
Earlier versions	Up to 12.16.1020	Driver will fail to load and show "Yellow Bang" in the device manager.		No limitations
1.50 and 1.60	Between 1x.16.1020 and 1x.19.2002 (IPoIB supported)	"Yellow Bang" unsupported mode - disable SR-IOV via mlxconfig	OK	No limitations
1.70 and onwards	1x.19.2002 and onwards (IPoIB supported)	OK	OK	No limitations

For further information on how to enable/disable SR-IOV, please refer to section [Single Root I/O Virtualization \(SR-IOV\)](#).