NVIDIA Native Drivers for VMware ESXi Inbox Drivers Release Notes

Exported on Nov/03/2023 09:12 AM

Table of Contents

Overview

These are the release notes of NVIDIA® ConnectX®-4 onwards adapter cards and NVIDIA® BlueField DPU ESXi Inbox Drivers which support the following uplink speeds:

Inbox Driver Version	OS	Uplink Speed
4.23.0.66-1vmw	ESXi 8.0u2	10/25/40/50/100/200GbE
4.23.0.36-12vmw	ESXi 8.0u1	10/25/40/50/100/200GbE
4.23.0.36-8vmw	ESXi 8.0	10/25/40/50/100/200GbE
4.17.16.10-vmw	ESXi 7.0u2	10/25/40/50/100/200GbE
4.17.16.8-vmw	ESXi 7.0u1	10/25/40/50/100GbE
4.17.16.7-vmw	ESXi 7.0	10/25/40/50/100GbE
4.17.13.1-vmw	ESXi 6.7u2	10/25/40/50/100GbE
4.17.9.12-vmw	ESXi 6.7	10/25/40/50/100GbE

General Support

NVIDIA BlueField-2 DPU for VMware vSphere Distributed Services Engine

The Inbox drivers support the following NVIDIA BlueField-2 DPU components:

	Recommended Version
Supported OSes	ESXi 8.0u2
NVIDIA BlueField-2 Firmware	24.36.7506
Arm Firmware (ATF/UEFI)	4.0.2.12722
BMC	23.04-3

The firmware and the driver are supplied by Server OEM for VMware vSphere Distributed Services Engine.

NVIDIA Ethernet HCAs

The Inbox drivers support the following Ethernet HCA:

HCAs	Recommended Firmware Rev.	Supported OSes	Additional Information
BlueField-3* (NIC Mode)	32.38.3056	ESXi 8.0u2	• DOCA: 2.2.1 • Arm Firmware (ATF/ UEFI): - 4.2.2 • BMC: 23.09-6
BlueField-2*	24.38.1002	ESXi 8.0u2, ESXi 8.0u1,ESXi 8.0	
ConnectX-7	28.38.1002	ESXi 8.0u2	
ConnectX-6 Dx	22.38.1002	ESXi 8.0u2, ESXi 8.0u1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1	
ConnectX-6 Lx	26.38.1002	ESXi 8.0u2, ESXi 8.0u1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1	
ConnectX-6	20.38.1002	ESXi 8.0u2, ESXi 8.0u1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0	
ConnectX-5 / ConnectX-5 Ex	16.35.2000	ESXi 8.0u2, ESXi 8.0u1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0, ESXi 6.7u2, ESXi 6.7	
ConnectX-4 Lx	14.32.1010		
ConnectX-4	12.28.2006		

* Non-VMware vSphere Distributed Services Engine (acting as ConnectX device)

Supported Features/Capabilities

⚠ New VMware OS releases include all features from previous releases unless otherwise specified in this document.

Feature/ Change	Description	DPU Only	Support Added in OS	Supported in Version
SR-IOV	Updated the default VF number in VMware vSphere Distributed Services Engine devices to 64.	Yes*	ESXi 8.0 U2	4.23.0.66-1vmw
General	Added support for up to 8 x 100G over DVS (UENS) interfaces.	No	ESXi 8.0 U2	4.23.0.66-1vmw
ENS	Enabled Netqueue RSS for ENS by default.	No	ESXi 8.0 U2	4.23.0.66-1vmw
Communication Channel	The communication channel between x86 architecture and the DPU is now exposed via a separate PCI device.	Yes*	ESXi 8.0 U2	4.23.0.66-1vmw
Adapter Cards	Added support for NVIDIA BlueField-3 DPU at Technical Preview level.	Yes*	ESXi 8.0 U2	4.23.0.66-1vmw
Adapter Cards	Added support for ConnectX-7 adapter cards.	No	ESXi 8.0 U2	4.23.0.66-1vmw
SR-IOV	Added support for ConnectX-7 SR-IOV IB.	No	ESXi 8.0 U2	4.23.0.66-1vmw
DPU Offloads Enable/Disable	DPU offloads can be now enabled/disabled using the "disable_dpu_accel" module parameter. The following are the new parameter's values: • 0 - Full acceleration (Default) • 1 - No acceleration, classic NIC mode **Note:** The same value must be used for both the host and the Arm side	Yes*	ESXi 8.0 / ESXi 8.0 U1	4.23.0.36-8vmw / 4.23.0.36-12vmw
NVIDIA Proprietary Flow Statistics	Statistical information on flows and actions that can be queried by the user using the nmlxcli tool. For further information, please refer to the User Manual and nmlxcli.	Yes*		
ENS VXLAN Encap/Decap Offload	VXLAN hardware encapsulation and decapsulation offload enables the traditional offloads to be performed on the encapsulated traffic. With NVIDIA BlueField-2 SmartNICs, data center operators can decouple the overlay network layer from the physical NIC performance, thus achieving native performance in the new network architecture. The VXLAN encapsulation and decapsulation offload in VMware is done using the NSX manager and the FPO interface which can issue FPO rules to instruct the NVIDIA hardware to perform the VXLAN encapsulation and decapsulation offload.	Yes*		

Feature/ Change	Description	DPU Only	Support Added in OS	Supported in Version
NSX Edge Acceleration	Added support for running NSX edge VM with UPT interfaces form factor on ESXi using accelerated DVS.	Yes*		
Hardware Accelerated LAG	Added support for hardware accelerated LAG. When two uplinks are added for a virtual switch, a hardware LAG is created between the two uplinks to provides connectivity between the Virtual Functions on different Physical Functions. Additionally, it allows teaming policies to be set on the LAG to facilitate teaming actions such as failover.	Yes*		
Hardware Accelerated Uniform Pass-Through (UPT)	UPT interface is provided by the DPU to the host accelerating data-path directly to the hardware.	Yes*		
ENS Flow Processing Offload (FPO) Model 3 Support	The driver can now support ENS FPO Model 3 (i.e. support SR-IOV with logical L2 and L3 offloaded to E-Switch and support UPT).	Yes*		
Reading Temperature Sensors	Enables the driver to read the temperature from private statistics. To see the temperature, run: `# nsxdp-cli ens uplink stats get -n vmnic1 | grep -i asicSensorTemperature`	No		
Classic NIC Mode	In embedded CPU mode, when no ENS is present, the default miss rule will forward all traffic to the x86 PF.	Yes*		
Hardware Accelerated RoCE Encapsulation	Added support for Hardware Accelerated GENEVE and VXLAN encapsulation and decapsulation for RoCE traffic.	No		
Hardware Accelerated Packet Capture	Hardware accelerated flows can now be mirrored using the standard packet capture tools.	No		
Hardware Accelerated NSX Distributed Firewall	Added the ability to offload NSX Distributed Firewall rules by using in-hardware tracking of packet flows.	No		
Receive Side Scaling (RSS) for ENS Model 0 and Model 1	RSS support for ENS Model 1 improves performance using fewer CPU cores. This capability can be enabled using the `"netq_rss_ens"` module parameter.	No		
vSAN over RDMA	vSAN over RoCE brings performance improvements for vSAN technology, offloads the CPU from performing the data communication tasks, significantly boosts overall storage access performance, and enables massive amount of data transfers.	No	ESXi 7.0u2	4.19.16.10-vmw
Adapter Cards	Added support for NVIDIA BlueField-2 adapter cards	Yes		

Feature/ Change	Description	DPU Only	Support Added in OS	Supported in Version
Software DCBx	Data Center Bridging (DCB) uses DCBX to exchange configuration information with directly connected peers. DCBX operations can be configured to set PFC or ETS values.	No	ESXi 7.0u1	4.19.16.8-vmw
RDMA Native Endpoint Support	Added support for RDMA communication with RDMA Native endpoints. RDMA Native endpoints are RDMA capable devices such as storage arrays that do not use the PVRDMA adaptor type (non-PVRDMA endpoints).	No		
Adapter Cards	Added support for NVIDIA ConnectX-6 Dx adapter cards	No		
Differentiated Services Code Point (DSCP)	DSCP is a mechanism used for classifying network traffic on IP networks. It uses the 6-bit Differentiated Services Field (DS or DSCP field) in the IP header for packet classification purposes. Using Layer 3 classification enables you to maintain the same classification semantics beyond local network, across routers. Every transmitted packet holds the information allowing network devices to map the packet to the appropriate 802.1Qbb CoS. For DSCP based PFC or ETS the packet is marked with a DSCP value in the Differentiated Services (DS) field of the IP header.	No	ESXi 7.0	4.19.16.7-vmw
Insufficient Power Detection	Enables the driver to report when Insufficient PCI power is detected.	No		
PVRDMA Namespace	Creates RDMA resources with a specified ID.	No		
Adapter Cards/ SmartNIC	Added support for NVIDIA ConnectX-6 and NVIDIA BlueField cards.	No		
Dynamic RSS	Improves network performance by allowing OS Load Balancer better RSS RX queue utilization during heavy traffic of the same type. For further information, see section "*Dynamic RSS*" in the User Manual.	No	ESXi 6.7u2	4.17.13.1-vmw
Packet Capture Utility (Sniffer)	Packet Capture utility duplicates all traffic, including RoCE, in its raw Ethernet form (before stripping) to a dedicated "sniffing" QP, and then passes it to an ESX drop capture point. It allows gathering of Ethernet and RoCE bidirectional traffic via pktcap-uw and viewing it using regular Ethernet tools, e.g. Wireshark. For further information, see section *"Packet Capture Utility"* in the User Manual.	No		

Feature/ Change	Description	DPU Only	Support Added in OS	Supported in Version
SR-IOV max_vfs module parameter Type Modification	Changed the type of the SR-IOV max_vfs module parameter from a single integer value to an array of unsigned integers. For further information, refer to the User Manual.	No	ESXi 6.7	4.17.9.12-vmw
DCBX Negotiation Support for PFC	PFC port configuration can now be auto-negotiated with switches that support the DCBX protocol.	No		
ESXi CLI	ESXi CLI support for ESXi 6.7	No		
Geneve Stateless Offload	Geneve network protocol is encapsulated into IP frame (L2 tunneling). Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header.	No		
Remote Direct Memory Access (RDMA)	Remote Direct Memory Access (RDMA) is the remote memory management capability that allows server-to-server data movement directly between application memory without any CPU involvement. **Note:** It is recommended to use RoCE with PFC enabled in driver and network switches. For how to enable PFC in the driver see section "*Priority Flow Control (PFC)*" in the User Manual.	No		
Set Link Speed	Enables you to set the link speed to a specific link speed supported by ESXi. For further information, see section *"Set Link Speed"* in the User Manual.	No		
Priority Flow Control (PFC)	Applies pause functionality to specific classes of traffic on the Ethernet link. For further information, see section *"Priority Flow Control (PFC)"* in the User Manual.	No		
NetQ RSS	Allows the user to configure multiple hardware queues backing up the sin- gle RX queue. NetQ RSS improves vMotion performance and multiple streams of IPv4/IPv6 TCP/UDP/IPSEC bandwidth over single interface between the Virtual Machines. For further information, see section *"NetQ RSS"* in *the User Manual*.	No		
Default Queue RSS (DRSS)	Allows the user to configure multiple hardware queues backing up the default RX queue. DRSS improves performance for large scale multicast traffic between hypervisors and Virtual Machines interfaces. For further information, see section *"Default Queue Receive Side Scaling (DRSS)"* in the User Manual.	No		
SR-IOV	Single Root IO Virtualization (SR-IOV) is a technology that allows a physical PCIe device to present itself multiple times through the PCIe bus.	No		

Feature/ Change	Description	DPU Only	Support Added in OS	Supported in Version
	Support for up to 8 ConnectX-4 ports and up to 16 VFs. For further information, refer to the User Manual	No		
RX/TX Ring Resize	Allows the network administrator to set new RX\TX ring buffer size.	No		
VXLAN Hardware Stateless Offloads	Added support for VXLAN hardware offload. VXLAN hardware offload enables the traditional offloads to be performed on the encapsulated traffic. With ConnectX®-3 Pro, data center operators can decouple the overlay network layer from the physical NIC performance, thus achieving native performance in the new network architecture.	No		
NetDump	Enables a host to transmit diagnostic information via the network to a remote netdump service, which stores it on disk. Network-based coredump collection can be configured in addition to or instead of disk-based core-dump collection.	No		
NetQueue	NetQueue is a performance technology in VMware ESXi that significantly improves performance in Ethernet virtualized environments.	No		
Wake-on-LAN (WoL)	Allows a network administrator to remotely power on a system or to wake it up from the sleep mode.	No		
Hardware Offload	• Large Send Offload (TCP Segmentation Offload) • RSS (Device RSS)	No		
Hardware Capabilities	• Multiple Tx/Rx rings • Fixed Pass-Through • Single/Dual port • MSI-X	No		
Ethernet Network	• TX/RX checksum • Auto moderation and Coalescing • VLAN stripping offload	No		

* Note: Supported only on VMware certified VMware vSphere Distributed Services Engine system.

RSS Additional Information

Driver	Default Queue RSS	Netqueue RSS	Device RSS
4.23.0.66-1vmw	Non ENS ENS Model0 ENS Model1 ENS Model3 (On host)	Non ENS ENS Model0 ENS Model1 ENS Model3 (On host)	Non ENS ENS Model0 ENS Model1 ENS Model3 (On host)
4.23.0.36-12vmw	Non ENS ENS Model0 ENS Model1 ENS Model3 (On host)	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default) ENS Model3 (On host)	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default) ENS Model3 (On host)

Driver	Default Queue RSS	Netqueue RSS	Device RSS
4.23.0.36-8vmw	Non ENS ENS Model0 ENS Model1 ENS Model3 (On host)	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default) ENS Model3 (On host)	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default) ENS Model3 (On host)
4.22.73.1004	Non ENS ENS Model0 ENS Model1	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default)	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default)
4.21.71.101	Non ENS ENS Model0 ENS Model1	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default)	Non ENS
4.21.71.1	Non ENS ENS Model0 ENS Model1	Non ENS ENS Model0 (disabled by default) ENS Model1 (disabled by default)	Non ENS
4.19.71.1	Non ENS	Non ENS	Non ENS
4.19.70.1	Non ENS	Non ENS	Non ENS
4.17.15.16	Non ENS	Non ENS	Non ENS

Known Issues

The following is a list of general limitations and known issues of the various components of the Inbox Driver release.

Internal Ref.	Issue
3432688	**Description:** The PCI device with device ID 0xc2d1 (BlueField Auxiliary Comm Channel) is used as the communication channel between the DPU and Host and is essential for SmartNIC operation. Therefore, it must not be enabled for passthrough.
	Workaround: N/A
	Keywords: BlueField, communication channel, DPU, Host, passthrough
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0 U2
	Discovered in Version: 4.23.0.66
3588655	**Description:** LRO is not supported in UPT.
	Workaround: N/A
	Keywords: UPT, LRO, Performance
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0
	Discovered in Version: 4.23.0.66
3343206	**Description:** Universal Pass Through (UPT) tunneling is supported only when using ESXi 8.0u2 with the following firmware driver version combinations. Other combination are not recommended and can lead to driver's issues. • Firmware: xx.36.1010 • Driver: 4.24.0.1-1vmw
	Workaround: N/A
	Keywords: UPT tunneling
	Adapter Cards / DPU: ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0
	Discovered in Version: 4.23.0.66
3442918	**Description:** There is no VLAN traffic when setting VLAN in the VM for a VF interface which performs PCI passthrough (not an SR-IOV passthru) to the VM.
	Workaround: N/A
	Keywords: VLAN
	Adapter Cards / DPU: ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.66
3136502	**Description:** In case of low/no memory in the system, the driver may get stuck in an endless loop trying to allocate memory.
	Workaround: N/A
	Keywords: System memory
	Adapter Cards / DPU: ConnectX-6 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.66

Internal Ref.	Issue
3436579	**Description:** Performance degradation might be experienced when running with LRO enabled in the VM due to a vmxnet3 driver bug.
	Workaround: To overcome the issue, perform the following: 1. Update to OS with kernel 6.3 and newer which contains the vmxnet3 driver fix: *3bced313b9a5 vmxnet3: use gro callback when UPT is enabled*. 2. Turn off LRO in the VM.
	Keywords: LRO, performance, vmxnet3 driver
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.66
3390706	**Description:** The maximum number of VFs for a PF on ConnectX-4 and ConnectX-5 adapter cards is 126 although the OS reports that the device supports 127 VFs.
	Workaround: Set the maximum number of VFs for a PF using the nmlx5_core module parameter max_vfs, or by using the ESXi management tools.
	Keywords: VF, PF, ConnectX-4, ConnectX-5
	Adapter Cards / DPU: ConnectX-4, ConnectX-5
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.66
3344305	**Description:** Enabling both UPT VFs and SR-IOV VFs on the same host will result in IOMMU fault.
	Workaround: N/A
	Keywords: UPT SR-IOV
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
2678029	**Description:** Due to hardware limitations, Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is **not** supported in **ConnectX-5 and ConnectX-6** adapter cards.
	Workaround: Use ConnectX-6 Lx, ConnectX-6 Dx, or onwards cards that support ENS Model 1 Level 2 and Model 2A.
	Keywords: ENS, ConnectX-5/ConnectX-6, Model 1 Level 2 and Model 2A
	Adapter Cards / DPU: ConnectX-4 Onwards HCAs
	Available in OS: ESXi 7.0u3, ESXi 7.0u2
	Discovered in Version: 4.23.0.36
-	**Description:** Geneve options length support is limited to 56B. Received packets with options length bigger than 56B are dropped.
	Workaround: N/A
	Keywords: Geneve
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
-	**Description:** The hardware can offload only up to 256B of headers.
	Workaround: N/A
	Keywords: Hardware offload

Internal Ref.	Issue
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2204581	**Description:** A mismatch between the uplink and the VF MTU values may result in CQE with error.
	Workaround: Align the uplink and the VF MTU values.
	Keywords: CQE, error, model 2,
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2429623	**Description:** Enabling sriov_mc_isolation module parameter may result in vmknic and emulated NICs multicast and IPv6 traffic loss.
	Workaround: Unset or set the module parameter to 0.
	Keywords: Multicast, IPv6, SR-IOV
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2372060	**Description:** RDMA is not supported in the Hypervisor with ENS model 2.
	Workaround: N/A
	Keywords: ENS model 2, RDMA
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2139469	**Description:** Setting the "Allow Guest MTU Change" option in vSphere Client is currently not functional. Although guest MTU changes in SR-IOV are allowed, they do not affect the port's MTU and the guest's MTU remains the same as the PF MTU.
	Workaround: N/A
	Keywords: MTU, SR-IOV
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
1340255	**Description:** ECN statistic counters accumulatorsPeriod and ecnMarkedRocePackets display wrong values and cannot be cleared.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36

Internal Ref.	Issue
1340275	**Description:** ECN tunable parameter initialAlphaValue for the Reaction Point protocol cannot be modified.
	Workaround: N/A
	Keywords: nmlx5 ecn nmlxcli
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2430662	**Description:** Card's speed remains zero after port goes down and reboot is performed.
	Workaround: Turn the port down and then up again.
	Keywords: ConnectX-6 Dx, link speed
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
1514289	**Description:** RoCE traffic may fail after vMotion when using namespace.
	Workaround: N/A
	Keywords: Namespace, RoCE, vMotion
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2334405	**Description:** Legacy SR-IOV is not supported with Model 1.
	Workaround: Unset max_vfs or alternatively move to ENS model 0 or Model 2.
	Keywords: SR-IOV, ENS
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2449578	**Description:** When in ENS mode, changing the scheduler to HCLK, may cause traffic loss.
	Workaround: N/A
	Keywords: ENS, HCLK scheduler
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
746100	**Description:** The 'esxcli mellanox uplink link info -u <vmnic_name>' command reports the 'Auto negotiation' capability always as 'true'.
	Workaround: N/A
	Keywords: 'Auto negotiation' capability
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0

Internal Ref.	Issue
	Discovered in Version: 4.23.0.36
1068621	**Description:** SMP MADs (ibnetdiscover, sminfo, iblinkinfo, smpdump, ibqueryerr, ibdiagnet and smpquery) are not supported on the VFs.
	Workaround: N/A
	Keywords: SMP MADs
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0 U2, ESXi 8.0 U1, ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
1446060	**Description:** Although the max_vfs module parameter range is "0-128", due to firmware limitations, the following are the supported VFs per single port devices: • ConnectX-4 / ConnectX-5: up to 127
	Workaround: N/A
	Keywords: SR-IOV, VFs per port
	Adapter Cards / DPU: ConnectX-4 Onwards HCAs
	Available in OS: ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in Version: 4.23.0.36
2139469	**Description:** Setting the "Allow Guest MTU Change" option in vSphere Client is currently not functional. Although guest MTU changes in SR-IOV are allowed, they do not affect the port's MTU and the guest's MTU remains the same as the PF MTU.
	Workaround: N/A
	Keywords: MTU, SR-IOV
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
2813820	**Description:** The maximum supported VFs in the case of Model 3 (A and B) setup is 126.
	Workaround: Set the maximum value of NUM_OF_VFS to 126 in mlxconfig.
	Keywords: VFs, Model 3
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
2615605	**Description:** LAG creation fails when VFs are enabled on the host.
	Workaround: Before adding a second uplink to the switch, make sure all VMs with VFs are powered off.
	Keywords: Link Aggregation
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
2611342	**Description:** On older version of iDRAC, the NC-SI channel is not opened automatically.
	Workaround: Run this command from the iDrac root SHEll manually and open the channel after every reboot or Arm OS `libncsitest 47 eth2 0 2 <MAC addr of sd0>  <MAC addr of sd0>`
	Keywords: iDRAC Communication

Internal Ref.	Issue
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
-	**Description:** Adapter cards that come with a pre-configured link type as InfiniBand cannot be detected by the driver and cannot be seen by MFT tools. Thus its link type cannot be changed.

	Workaround:
	1. Unload the driver.
	`unload nmlx5_core module`
	2. Make the device visible to MFT by loading the driver in a recovery mode.
	`vmkload_mod nmlx5_core mst_recovery=1 kill the devmgr`
	3. Check the device available on your machine.
	`/opt/mellanox/bin/mst status`
	4. Change the link type to Ethernet using MFT.
	`opt/mellanox/bin/mlxconfig -d mt4115_pciconf0 set LINK_TYPE_P1=2 LINK_TYPE_P2=2`
	5. Power Cycle the host.

	Keywords: Link type, InfiniBand, MFT
	Adapter Cards / DPU: ConnectX-4 Onwards HCAs
	Available in OS: ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0
	Discovered in versions: 4.17.13.10-vmw, 4.17.13.1-vmw, 4.17.16.8-vmw, 4.17.16.7-vmw, 4.17.9.12-vmw

Bug Fixes

Internal Ref.	Issue
3175465	**Description:** The driver version in iDRAC is not displayed correctly.
	Keywords: iDRAC
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine, ConnectX-4 Onwards HCAs
	Available in OS: ESXi 8.0, ESXi 7.0u3, ESXi 7.0u2, ESXi 7.0u1, ESXi 7.0, ESXi 6.7u2, ESXi 6.7
	Discovered in Version: 4.23.0.36
	Fixed in Release: 4.23.0.66
3038613	**Description:** On DPU, the following error message can appear upon driver loading: *"ECPF is not EswManager, cannot get SF support on host"*. This message is harmless and can be ignored.
	Keywords: Error, log
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
	Fixed in Release: 4.23.0.66
3192250	**Description:** The following error in logs may appear when performing DFW HW offload: *"HW state is LAST_ACK while both sides sent an ACK"*. This message is harmless and can be ignored.
	Keywords: DFW, offload, DPU
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
	Fixed in Release: 4.23.0.66

3308850	**Description:** To create 126 UPT devices, the following configuration needs to be set manually:
	`# /opt/mellanox/bin/mlxconfig -d <dev> s NUM_OF_VFS=64`
	`# /opt/mellanox/bin/mlxconfig -d <dev> s NUM_PF_MSIX=127`

	Keywords: UPT
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
	Fixed in Release: 4.23.0.66
-	**Description:** Enabling VMDirectPath I/O pass-through of the first interface is not supported.
	Keywords: Pass-through
	Adapter Cards / DPU: NVIDIA BlueField-2 for VMware vSphere Distributed Services Engine
	Available in OS: ESXi 8.0
	Discovered in Version: 4.23.0.36
	Fixed in Release: 4.23.0.66
3109870	**Description:** Fixed an issue that caused the driver configuring the wrong SL value for RoCE with RDMACM.
	Keywords: RoCE, QOS
	Fixed in Release: 4.23.0.36

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: “NVIDIA”) make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.