



ConnectX[®]-6 Dx Ethernet Network Interface Card IC

Intelligent & Secure Cloud Computing

Industry's most intelligent and secure cloud Ethernet network adapter (NIC) ASIC, accelerating mission-critical cloud applications, e.g., security, virtualization, SDN/NFV, big data, machine learning, and storage, and delivering the highest return on investment (ROI)

ConnectX-6 Dx is a member of the award-winning ConnectX series of network adapters whose novel capabilities accelerate cloud and data-center payloads with unprecedented world-class performance and resiliency. ConnectX-6 Dx adapters are equipped with either dual ports of 10/25/40/50/100 GbE or a single port of 200GbE, delivering sub-800ns latency and up to 215 million messages per second.

Advanced Virtualization

ConnectX-6 Dx delivers another level of innovation to enable the building of highly efficient virtualized cloud data centers:

Virtualization – Mellanox ASAP² - Accelerated Switch and Packet Processing[®] technology for vSwitch/vRouter hardware offload, delivers orders of magnitude higher performance vs. software-based solutions. ConnectX-6 Dx ASAP² offers both SR-IOV and VirtIO in-hardware offload capabilities, and supports up to 8 million rules.

Advanced Quality of Service – Includes traffic shaping, and classification-based data policing.

Security From Zero Trust to Hero Trust

In an era where privacy of information is key and zero trust is the rule, ConnectX-6 Dx adapters offer a range of advanced built-in capabilities that bring security down to the endpoints with unprecedented performance and scalability, including:

Crypto – IPsec and TLS data-in-motion inline encryption and decryption offload and AES-XTS block-level data-at-rest encryption and decryption offload.

Probes & DoS Attacks Protection – ConnectX-6 Dx enables a hardware-based L4 firewall by ASAP² offloading of stateful connection tracking.

NIC Security – Hardware Root-of-Trust (RoT) Secure Boot and secure firmware update using RSA cryptography, and cloning-protection, via a device-unique secret key.

Industry-leading RoCE

Following the Mellanox ConnectX tradition of industry-leading RoCE capabilities, ConnectX-6 Dx adds another layer of innovation to enable more scalable, resilient and easy-to-deploy RoCE solutions.

Zero Touch RoCE – ConnectX-6 Dx allows RoCE payloads to run seamlessly on existing networks without requiring special network configuration (no PFC, no ECN) for simplified RoCE deployments. ConnectX-6 Dx's new features ensure resiliency and efficiency at scale of such deployments.

Configurable Congestion Control – API to build user-defined congestion control algorithms, best serving various environments and RoCE and TCP/IP traffic patterns.

Machine Learning and Big Data

Machine learning applications are based on training a deep neural network, which requires complex computations and fast and efficient data delivery. Mellanox solutions enable smart offloading such as RDMA, GPUDirect[®] and more advanced capabilities that dramatically improve neural network training performance and overall machine learning applications. GPUDirect was developed to dramatically improve GPU-to-GPU communication, reducing latency and increasing performance and CPU utilization.

Efficient Storage Solutions

The evolving NVMe over Fabric (NVMe-oF) protocol leverages RDMA connectivity to remotely access NVMe storage devices efficiently, while keeping the end-to-end NVMe model at lowest latency. With its NVMe-oF target and initiator offloads, ConnectX-6 Dx brings further optimization to NVMe-oF, enhancing CPU utilization and scalability. ConnectX-6 Dx supports hardware offload for ingress/egress of T10-DIF/PI/CRC32/CRC64 signatures. Additionally, ConnectX-6 Dx enables Host Chaining, an innovative storage rack design by which different servers can be connected without a switch.

ConnectX[®]-6 Dx

HIGHLIGHTS

Features & Applications

- 10/25/40/50/100/200 Gb/s Ethernet, PAM4/NRZ
- Up to 200Gb/s bandwidth
- Message rate of up to 215Mpps
- Sub 0.8usec latency
- Mellanox Multi-Host with advanced QoS
- ASAP² - Accelerated Switching and Packet Processing for virtual switches/routers
- Supports a variety of overlay tunnels
- Network service chaining acceleration for SDN and NFV
- Stateful rule checking for connection tracking
- IPsec and TLS in-line crypto acceleration
- Block crypto acceleration for data-at-rest
- Hardware Root-of-Trust and secure firmware update
- Advanced RoCE capabilities
- Best in class PTP for TSN applications
- Host chaining technology for economical rack design
- Platform agnostic: x86, Power, Arm

Solutions

- Enterprise data-centers
- Cloud-native, Web 2.0, hyperscale
- Cyber security
- Big data analytics
- Scale-out compute and storage infrastructure
- Telecom and Network Function Virtualization (NFV)
- Cloud storage
- Machine Learning (ML) and Artificial Intelligence (AI)
- Low latency, high frequency financial trading
- High-performance computing (HPC)
- Media and Entertainment



Features*

Virtualization

- VirtIO in-hardware
- Single Root IOV (SR-IOV)
 - 1K VFs per port
 - 64 PFs (16 per host)
- VMware NetQueue support
- Per VM QoS
- Hardware offloads
 - Connection tracking
 - Header rewrite
 - Stateless offloads for overlay network tunneling protocols
 - Encapsulation and decapsulation of VXLAN, NVGRE, GENEVE and more overlay networks

Cyber Security

- IPsec encryption & decryption
 - AES-GCM 128/256 bit key
- TLS encryption & decryption
 - AES-GCM 128/256 bit key
- AES-XTS data-at-rest encryption & decryption
- AES-GCM 128/256 bit key
- Hardware root-of-trust
- Secure firmware update

Networking Offloads

- TCP/UDP/IP stateless offload
- LSO, LRO, checksum offload
- RSS (also on encapsulated packet), TSS, VLAN and MPLS tag insertion/stripping, receive flow steering

Networking Accelerations

- Data Plane Development Kit (DPDK)
- Hardware-based and software-enabled XDP acceleration

Storage

- NVMe over Fabric offloads for target
- Storage protocols: SRP, iSER, NFSoRDMA, SMB Direct, NVMe-of, and more
- T10 DIF - signature handover
- CRC16, CRC32 and CRC64 signature offloads

Time Seventieths Networking (PTP)

- Clock synchronization accuracy better than 16ns
- Software controlled hardware clock
- Full wire-speed hardware time stamping
- PPS in/out with programmable frequency
- High precision oscillator options
- Accurate package scheduling
- IEEE 1588v2 (PTP)
 - Supports OC, SC, BC, MC PTP Clocks
 - One & two step sync methods
 - E2E and P2P
 - Integrated with any PTP daemon

RDMA over Converged Ethernet (RoCE)

- Zero-touch network configuration
- Selective repeat
- Programmable congestion control mechanism

- Mellanox PeerDirect® RDMA (aka GPUDirect®)
- Dynamically Connected transport (DCT)
- Burst buffer offload
- RoCE over overlay networks

Mellanox Multi-Host

- Independent PCIe interfaces to independent hosts
- Two PCIe x8 to two hosts, or four PCIe x4 to four hosts
- Multi-Host congestion handling
- Independent SMBus interfaces
- Independent stand-by and wake-on-LAN signals
- Supports ECN marking capability in hardware

Management & Control

- NC-SI, MCTP over SMBus and MCTP over PCIe - Baseboard Management Controller interface
- PLDM for Monitor and Control DSP0248
- PLDM for Firmware Update DSP026
- I²C interface for device control and configuration
- General Purpose I/O pins
- SPI interface to flash
- JTAG IEEE 1149.1 and IEEE 1149.6

Media & Entertainment

- SMPTE 2110 streaming offload
 - 2110-10, 20, 21N, 30, 40, 50
 - ST2022-6/7
 - SMPTE ST-2059-2 PTP profile
 - Windows and Linux OS support

Remote Boot

- Remote boot over Ethernet
- Remote boot over iSCSI
- UEFI support
- PXE Boot

Host Interface

- PCIe Gen 4.0, 3.0, 2.0, 1.1
- 2.5, 5.0, 8, 16GT/s link rate
- 16 lanes of PCIe
- Support for PCIe x1, x2, x4, x8, and x16 configurations
- Socket Direct – overcoming QPI bottlenecks
- PCIe Atomic
- TLP (Transaction Layer Packet) Processing Hints (TPH)
- Embedded PCIe switch
- Advanced Error Reporting (AER)
- PCIe switch Downstream Port Containment (DPC) enablement for PCIe hot-plug
- Access Control Service (ACS) for peer-to-peer secure communication
- Process Address Space ID (PASID) Address Translation Services (ATS)
- IBM CAPIv2 (Coherent Accelerator Processor Interface)
- Support for MSI/MSI-X mechanisms

* This section describes hardware features and capabilities. Please refer to the driver and firmware release notes for feature availability.

Standards

- IEEE 802.3bs, 200 Gb/s
- IEEE 802.3cd, 50,100 and 200 Gb/s
- IEEE 802.3bj, 802.3bm 100 Gb/s
- IEEE 802.3by, 25, 50 Gb/s supporting all FEC modes
- IEEE 802.3ba 40 Gigabit Ethernet
- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE 802.3az Energy Efficient Ethernet (supports only "Fast-Wake" mode)
- IEEE 802.3ap based auto-negotiation and KR startup
- IEEE 802.3ad, 802.1AX Link Aggregation
- IEEE 802.1Q, 802.1P VLAN tags and priority
- IEEE 802.1Qaz (ETS)
- IEEE 802.1Qbb (PFC)
- IEEE 802.1Qbg
- 25G/50G Ethernet Consortium "Low Latency FEC" for 50/100/200GE PAM4 links
- PCI Express Gen 3.0 and 4.0

Table 1 - Part Numbers and Descriptions

Network Interface(s)	PCI Express Interface	Additional Features Supported	OPN
2x 100GbE	PCIe 3.0/4.0 x16	--	MT28928A0-NCCF-CE
2x 100GbE	PCIe 3.0/4.0 x16	Crypto enabled	MT28928A0-CCCF-CE
2x 100GbE	PCIe 3.0/4.0 x16	Mellanox Multi-Host	MT28928A0-NCCF-CEM
2x 100GbE	PCIe 3.0/4.0 x16	Mellanox Multi-Host, Crypto enabled	MT28928A0-CCCF-CEM
1x 200GbE	PCIe 3.0/4.0 x16	--	MT28924A0-NCCF-VE
1x 200GbE	PCIe 3.0/4.0 x16	Crypto enabled	MT28924A0-CCCF-VE
1x 200GbE	PCIe 3.0/4.0 x16	Mellanox Multi-Host	MT28924A0-NCCF-VEM
1x 200GbE	PCIe 3.0/4.0 x16	Mellanox Multi-Host, Crypto enabled	MT28924A0-CCCF-VEM