



NVIDIA MELLANOX CONNECTX-6 DX 100G ETHERNET SMARTNIC IC

NVIDIA® Mellanox® ConnectX®-6 Dx SmartNIC IC is the industry's most intelligent and secure cloud Ethernet network adapter (NIC) ASIC, accelerating mission-critical cloud applications, such as 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 200 GbE, delivering sub-800 ns latency and up to 215 million messages per second.

ADVANCED VIRTUALIZATION

ConnectX-6 Dx delivers another level of innovation to enable building 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.

Features & Applications

- > 10/25/40/50/100/200 Gb/s Ethernet, PAM4/NRZ
- > Up to 200 Gb/s bandwidth
- > Message rate of up to 215 Mpps
- > Sub 0.8 usec 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

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 – Simplifying RoCE deployments, ConnectX-6 Dx allows RoCE payloads to run seamlessly on existing networks without requiring special configuration on the network (no PFC, no ECN). New features in ConnectX-6 Dx 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.

Solutions

- > Enterprise data-centers
- > Cloud-native, Web 2.0, hyperscale
- > Cyber security
- > Big data analytics
- > Scale-out compute and storage infrastructure
- > Telco 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*

Network Interface

- > 2 x 25/50/100 GbE; 1 x 200 GbE

Host Interface

- > PCIe Gen 4.0, 3.0, 2.0, 1.1
- > 16.0, 8.0, 5.0, 2.5 GT/s link rate
- > 16 lanes of PCIe
- > MSI/MSI-X mechanisms
- > Advanced PCIe capabilities

Virtualization/Cloud Native

- > Single Root IOV (SR-IOV) and VirtIO acceleration
 - > Up to 1 K VFs per port
 - > 8 PFs
- > Support for tunneling
 - > Encap/decap of VXLAN, NVGRE, Geneve, and more
 - > Stateless offloads for Overlay tunnels

Mellanox ASAP²

- > SDN acceleration for:
 - > Bare metal
 - > Virtualization
 - > Containers
- > Full hardware offload for OVS data plane
- > Flow update through RTE_Flow or TC_Flower
- > OpenStack support
- > Kubernetes support
- > Rich classification engine (L2 to L4)
- > Flex-Parser: user defined classification
- > Hardware offload for:
 - > Connection tracking (L4 firewall)
 - > NAT
 - > Header rewrite
 - > Mirroring
 - > Sampling
 - > Flow aging
 - > Hierarchical QoS
 - > Flow-based statistics

Cyber Security

- > Inline hardware IPsec encryption and decryption
 - > AES-GCM 128/256 bit key
 - > IPsec over RoCE
- > Inline hardware TLS encryption and decryption
 - > AES-GCM 128/256 bit key
- > Data-at-rest AES-XTS encryption and decryption
 - > AES-GCM 128/256 bit key
- > Platform security
 - > Hardware root-of-trust
 - > Secure firmware update

Stateless Offloads

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

Advanced Timing & Synchronization

- > Advanced PTP
 - > IEEE 1588v2 (any profile)
 - > PTP Hardware Clock (PHC) (UTC format)
 - > 16 nsec accuracy
 - > Line rate hardware timestamp (UTC format)
 - > PPS In and configurable PPS Out
- > Time triggered scheduling
- > PTP based packet pacing
- > Time based SDN acceleration (ASAP2)
- > Time Sensitive Networking (TSN)

Storage Accelerations

- > NVMe over Fabric offloads for target
- > Storage protocols: iSER, NFSoRDMA, SMB Direct, NVMe-oF, and more
- > T-10 Dif/Signature Handover

RDMA over Converged Ethernet (RoCE)

- > RoCE v1/v2
- > Zero Touch RoCE: no ECN, no PFC
- > RoCE over overlay networks
- > IPsec over RoCE
- > Selective repeat
- > Programmable congestion control interface
- > GPUDirect[®]
- > Dynamically connected transport (DCT)
- > Burst buffer offload

Management and 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

Remote Boot

- > Remote boot over Ethernet
- > Remote boot over iSCSI
- > UEFI support for x86 and Arm servers
- > PXE boot

* 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

ORDERING INFORMATION

IC Network Interface	PCI Express Interface	Additional Features Supported	OPN
2 x 50GbE	PCIe 3.0/4.0 x16	-	MT28928A0-NCCF-GE
2 x 100GbE	PCIe 3.0/4.0 x16	-	MT28928A0-NCCF-CE
		Crypto enabled	MT28928A0-CCCF-CE
		Mellanox Multi-Host	MT28928A0-NCCF-CEM
		Crypto enabled, Mellanox Multi-Host or Socket Direct	MT28928A0-CCCF-CEM
1 x 200GbE	PCIe 3.0/4.0 x16	-	MT28924A0-NCCF-VE
		Crypto enabled	MT28924A0-CCCF-VE
		Mellanox Multi-Host	MT28924A0-NCCF-VEM
		Crypto enabled, Mellanox Multi-Host or Socket Direct	MT28924A0-CCCF-VEM

Only Crypto enabled devices can perform IPSec/TLS/AES_XTS encryption and decryption offload.
 All ConnectX-6 Dx IC devices support optional secure boot; contact Mellanox support for information on secure boot enablement and provisioning processes.

* This section describes hardware features and capabilities.

Please refer to the driver and firmware release notes for feature availability.

Learn more at www.mellanox.com/products/ethernet-adapter-ic/connectx-6-dx-ic

© 2020 Mellanox Technologies. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, ConnectX, Mellanox Multi-Host, Mellanox Socket Direct, GPUDirect, and ASAP² - Accelerated Switch and Packet Processing are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation 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. SEPT20/60286PB-R2

