

## NVIDIA MELLANOX CONNECTX-6 DX ETHERNET SMARTNIC

NVIDIA® Mellanox® ConnectX®-6 Dx SmartNIC is the industry's most secure and advanced cloud network interface card to accelerate mission-critical data-center applications, such as security, virtualization, SDN/NFV, big data, machine learning, and storage. The SmartNIC provides up to two ports of 100 Gb/s or a single-port of 200 Gb/s Ethernet connectivity and delivers the highest return on investment (ROI) of any smart network interface card.

ConnectX-6 Dx is a member of NVIDIA Mellanox's world-class, award-winning ConnectX series of network adapters powered by leading 50 Gb/s (PAM4) and 25/10 Gb/s (NRZ) SerDes technology and novel capabilities that accelerate cloud and data-center payloads.

### 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 Attack Protection – ConnectX-6 Dx enables a hardware-based L4 firewall by offloading stateful connection tracking through Mellanox ASAP<sup>2</sup> - Accelerated Switch and Packet Processing®.
- > 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.

### ADVANCED VIRTUALIZATION

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

- > Virtualization – Mellanox ASAP<sup>2</sup> technology for vSwitch/vRouter hardware offload delivers orders of magnitude higher performance vs. software-based solutions. ConnectX-6 Dx ASAP<sup>2</sup> 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.

### SmartNIC Portfolio

- > 1/10/25/40/50/100/200 Gb/s Ethernet, PAM4/NRZ
- > Various form factors:
  - > PCIe low-profile
  - > OCP 3.0 Small Form Factor (SFF)
  - > OCP 2.0
- > Connectivity options:
  - > SFP28, SFP56, QSFP28, QSFP56, DSFP
  - > PCIe Gen 3.0/4.0 x16 host interface
  - > Multi-host and single-host flavors
  - > Crypto and non-crypto versions

### Key Features

- > Up to 200 Gb/s bandwidth
- > Message rate of up to 215 Mpps
- > Sub 0.8 usec latency
- > Flexible programmable pipeline for new network flows
- > Mellanox Multi-Host with advanced QoS
- > ASAP<sup>2</sup> - Accelerated Switching and Packet Processing for virtual switches/routers
- > Overlay tunneling technologies
- > IPsec and TLS in-line crypto acceleration
- > Block crypto acceleration for data-at-rest
- > Hardware Root-of-Trust and secure firmware update
- > Connection Tracking offload
- > Advanced RoCE capabilities
- > Best in class PTP for TSN applications
- > GPUDirect® for GPU-to-GPU communication
- > Host chaining technology for economical rack design
- > Platform agnostic: x86, Power, Arm
- > ODCC compatible

## 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.

## BEST-IN-CLASS PTP FOR TIME SENSITIVE APPLICATIONS

Mellanox offers a full IEEE 1588v2 PTP software solution as well as time sensitive related features called 5T45G. Mellanox PTP and 5T45G software solutions are designed to meet the most demanding PTP profiles. ConnectX-6 Dx incorporates an integrated Hardware Clock (PHC) that allows the device to achieve sub-20 usec accuracy while offering various timing related functions, including time-triggered scheduling or time-based SND accelerations (time based ASAP<sup>2</sup>). Furthermore, 5T45G technology enables software applications to transmit front-haul (ORAN) compatible in high bandwidth. The PTP solution supports slave clock, master clock, and boundary clock.

Selected ConnectX-6 Dx SmartNICs provide PPS-Out or PPS-In signals from designated SMA connectors.

## EFFICIENT STORAGE SOLUTIONS

With its NVMe-oF target and initiator offloads, ConnectX-6 Dx brings further optimization to NVMe-oF, enhancing CPU utilization and scalability. Additionally, ConnectX-6 Dx supports hardware offload for ingress/egress of T10-DIF/PI/CRC32/CRC64 signatures, as well as AES-XTS encryption/decryption offload enabling user-based key management and a one-time-FIPS-certification approach.

## WIDE SELECTION OF NICs

ConnectX-6 Dx SmartNICs are available in several form factors including low-profile PCIe, OCP2.0 and OCP3.0 cards, with various network connector types (SFP28/56, QSFP28/56, or DSFP). The ConnectX-6 Dx portfolio also provides options for Mellanox Multi-Host<sup>®</sup> and Mellanox Socket Direct<sup>®</sup> configurations.

Mellanox Multi-Host<sup>®</sup> connects multiple compute or storage hosts to a single interconnect adapter and enables designing and building new scale-out compute and storage racks. This enables better power and performance management, while reducing capital and operational expenses.

Mellanox Socket Direct<sup>®</sup> technology brings improved performance to multi-socket servers, by enabling each CPU in a multi-socket server to directly connect to the network through its dedicated PCIe interface. This enables data to bypass the QPI (UPI) and the other CPU, improving latency, performance and CPU utilization.

## Solutions

- > Cloud-native, Web 2.0, hyperscale
- > Enterprise data-centers
- > Cyber security
- > Big data analytics
- > Scale-out compute and storage infrastructure
- > Telco and Network Function Virtualization (NFV)
- > Cloud storage
- > Machine Learning (ML) & Artificial Intelligence (AI)
- > 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 ASAP<sup>2</sup>

- > 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 & decryption
  - > AES-GCM 128/256 bit key
  - > IPsec over RoCE
- > Inline hardware TLS encryption & decryption
  - > AES-GCM 128/256 bit key
- > Data-at-rest AES-XTS encryption & 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 (ASAP<sup>2</sup>)
- > 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<sup>®</sup>
- > Dynamically connected transport (DCT)
- > Burst buffer offload

### Management and Control

- > NC-SI, MCTP over SMBus and MCTP over PCIe
  - Baseboard Management Controller interface, NCSI over RBT in OCP 2.0/3.0 cards
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP026
- > I<sup>2</sup>C interface for device control and configuration

### Remote Boot

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

## STANDARDS\*


- > IEEE 802.3bs, 200 Gigabit Ethernet
- > IEEE 802.3cd, 50, 100 and 200 Gigabit Ethernet
- > IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet
- > IEEE 802.3by, 25, 50 Gigabit Ethernet 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
- > 25/50 Ethernet Consortium "Low Latency FEC" for 50GE/100GE/200GE PAM4 links
- > PCI Express Gen 3.0 and 4.0

\* This section describes hardware features and capabilities. Please refer to the driver and firmware release notes for feature availability. When using Mellanox Socket Direct or Mellanox Multi-Host in virtualization or dual-port use cases, some restrictions may apply. For further details, contact Mellanox Customer Support.

# SMARTNIC PORTFOLIO & ORDERING INFORMATION

## PCIE HHHL FORM FACTOR

Max Network Speed	Interface Type	Supported Ethernet Speeds [GbE]	Host Interface [PCIe]	Ordering Part Number (OPN)		
				No Crypto, No Secure Boot	With Crypto <sup>1</sup> , No Secure Boot	With Crypto <sup>1</sup> , With Secure Boot
2 x 25 GbE	SFP28	1/10/25	Gen 4.0 x8	MCX621102AN-ADAT	MCX621102AE-ADAT	MCX621102AC-ADAT
			Gen 4.0 x16	MCX623102AN-ADAT	Contact Mellanox	MCX623102AC-ADAT
2 x 50 GbE	SFP56	1/10/25/50	Gen 4.0 x16	MCX623102AN-GDAT	MCX623102AE-GDAT	MCX623102AC-GDAT
1 x 100 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup>	Gen 4.0 x16	MCX623105AN-CDAT	Contact Mellanox	Contact Mellanox
2 x 100 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup>	Gen 4.0 x16	MCX623106AN-CDAT	MCX623106AE-CDAT	MCX623106AC-CDAT
	QSFP56 + PPS In/Out SMAs	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup>	Gen 4.0 x16	MCX623106PN-CDAT	MCX623106PE-CDAT	MCX623106PC-CDAT
	DSFP	1/10/25/50 <sup>2</sup> /100	Gen 4.0 x16	Contact Mellanox	Contact Mellanox	Contact Mellanox
1 x 200 GbE	QSFP56	10/25/40/50 <sup>2</sup> /100 <sup>3</sup> /200	Gen 4.0 x16	MCX623105AN-VDAT	MCX623105AE-VDAT	MCX623105AC-VDAT




1 Use Crypto enabled cards to utilize IPsec/TLS/AES-XTS encryption/decryption hardware offload.  
 2 50G can be supported as either 2x25G NRZ or 1x50G PAM4 when using QSFP56.  
 3 100G can be supported as either 4x25G NRZ or 2x50G PAM4 when using QSFP56.

By default, the above products are shipped with a tall bracket mounted; a short bracket is included as an accessory.

## OCP 3.0 SMALL FORM FACTOR

Max Network Speed	Interface Type	Supported Ethernet Speeds [GbE]	Host Interface [PCIe]	Ordering Part Number (OPN)		
				No Crypto, No Secure Boot	With Crypto <sup>1</sup> , No Secure Boot	With Crypto <sup>1</sup> , With Secure Boot
2 x 25 GbE	SFP28	1/10/25	Gen 4.0 x16	MCX623432AN-ADAB	MCX623432AE-ADAB	MCX623432AC-ADAB
2 x 50 GbE	SFP56	1/10/25/50	Gen 4.0 x16	MCX623432AN-GDAB	Contact Mellanox	MCX623432AC-GDAB
2 x 100 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup>	Gen 4.0 x16	MCX623436AN-CDAB	MCX623436AE-CDAB	MCX623436AC-CDAB
1 x 200 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup> /200	Gen 4.0 x16	MCX623435AN-VDAB	Contact Mellanox	MCX623435AC-VDAB




1 Use Crypto enabled cards to utilize IPsec/TLS/AES-XTS encryption/decryption hardware offload.  
 2 50G can be supported as either 2x25G NRZ or 1x50G PAM4 when using QSFP56.  
 3 100G can be supported as either 4x25G NRZ or 2x50G PAM4 when using QSFP56.

These OPNs are Single Host; contact Mellanox for OCP OPNs with Mellanox Multi-Host support.  
 The above OCP3.0 OPNs come with Thumbscrew (pull tab) brackets; contact Mellanox for additional bracket types, such as Internal Lock or Ejector latch.

## OCP 2.0 FORM FACTOR

Max Network Speed	Interface Type	Supported Ethernet Speeds [GbE]	Host Interface [PCIe]	Ordering Part Number (OPN)		
				No Crypto, No Secure Boot	With Crypto <sup>1</sup> , No Secure Boot	With Crypto <sup>1</sup> , With Secure Boot
1 x 100 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup>	Gen 4.0 x16	MCX623405AN-CDAN	Contact Mellanox	MCX623405AC-CDAN
1 x 200 GbE	QSFP56	1/10/25/40/50 <sup>2</sup> /100 <sup>3</sup> /200	Gen 4.0 x16	Contact Mellanox	Contact Mellanox	Contact Mellanox



1 Use Crypto enabled cards to utilize IPsec/TLS/AES-XTS encryption/decryption hardware offload.  
 2 50G can be supported as either 2x25G NRZ or 1x50G PAM4 when using QSFP56.  
 3 100G can be supported as either 4x25G NRZ or 2x50G PAM4 when using QSFP56.

These OPNs are Single Host; contact Mellanox for OCP OPNs with Mellanox Multi-Host or Mellanox Socket Direct support.  
 ConnectX-6 Dx PCIe 4.0 cards are backward compatible.

\*For illustration only. Actual products may vary.

Learn more at [www.mellanox.com/products/ethernet/connectx-smartnic](http://www.mellanox.com/products/ethernet/connectx-smartnic)

© 2020 Mellanox Technologies. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, ConnectX, Mellanox Multi-Host, Mellanox Socket Direct, GPU Direct, Mellanox PeerDirect, and ASAP<sup>2</sup> - 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. AUG20/60259PB-R3

