

NVIDIA CONNECTX-6 DX ETHERNET SMARTNIC

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

ConnectX-6 Dx is powered by leading 50Gb/s (PAM4) and 25/10 Gb/s (NRZ) SerDes technology and novel capabilities that accelerate cloud and data center payloads.

ZERO-TRUST SECURITY

In an era where data privacy is key, ConnectX-6 Dx adapters offer advanced, built-in capabilities that bring security down to the endpoints with unprecedented performance and scalability:

- > Crypto—IPsec and TLS data-in-motion inline encryption and decryption offload, and AES-XTS block-level, data-at-rest encryption and decryption offloads
- > Probes and denial-of-service (DoS) attack protection—ConnectX-6 Dx enables a hardware-based L4 firewall by offloading stateful connection tracking through NVIDIA ASAP² - Accelerated Switch and Packet Processing[®] offload technology
- > 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 enables building highly efficient virtualized cloud data centers:

- > Virtualization—ASAP² delivers virtual switch (vSwitch) and virtual router (vRouter) hardware offloads at orders-of-magnitude higher performance than 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 (QoS)—ConnectX-6 Dx 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 Gen3 and Gen4 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 215Mpps
- > Sub 0.8usec latency
- > Programmable pipeline for new network flows
- > NVIDIA Multi-Host with advanced QoS
- > ASAP² for vSwitches and vRouters
- > 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 time-sensitive networking (TSN) applications
- > NVIDIA GPUDirect[®] for GPU-to-GPU communication
- > Host chaining technology for economical rack design
- > Platform agnostic: x86, Power, Arm
- > Open Data Center Committee (ODCC) compatible

INDUSTRY-LEADING ROCE

With industry-leading capabilities, ConnectX-6 Dx delivers more scalable, resilient, and easy-to-deploy remote direct-memory access over converged Ethernet (RoCE) solutions.

- > Zero Touch RoCE (ZTR)—Simplifying RoCE deployments, ConnectX-6 Dx with ZTR allows RoCE payloads to run seamlessly on existing networks without special configuration, either to priority flow control (PFC) or explicit congestion notification (ECN). ConnectX-6 Dx ensures the resilience, efficiency, and scalability of deployments.
- > Programmable congestion control—ConnectX-6 Dx includes an API for building user-defined congestion control algorithms for various environments running RoCE and background TCP/IP traffic concurrently.

BEST-IN-CLASS PTP FOR TIME SENSITIVE APPLICATIONS

NVIDIA offers a full IEEE 1588v2 Precision Time Protocol (PTP) software solution as well as time-sensitive-related features called 5T for 5G. NVIDIA PTP and 5T for 5G software solutions are designed to meet the most demanding PTP profiles. ConnectX-6 Dx incorporates an integrated PTP hardware clock (PHC) that allows the device to achieve sub-20 nanosecond (nsec) accuracy while offering timing-related functions, including time-triggered scheduling or time-based, software-defined networking (SDN) accelerations (time based ASAP²). 5T for 5G technology also enables software applications to transmit front-haul radio area network (RAN)-compatible data in high bandwidth. The PTP solution supports slave clock, master clock, and boundary clock operations.

ConnectX-6 Dx also supports SyncE, allowing selected ConnectX-6 Dx SmartNICs to 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, enhancing CPU utilization and scalability. Additionally, ConnectX-6 Dx supports hardware offload for ingress and egress of T10-DIF/PI/CRC32/CRC64 signatures and AES-XTS encryption and decryption offloads, enabling user-based key management and a one-time Federal Information Processing Standards (FIPS) certification approach.

NIC PORTFOLIO

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 NVIDIA Multi-Host[®] and NVIDIA Socket Direct[®] configurations.

ConnectX-6 Dx adds significant improvements to NVIDIA Multi-Host applications by offering advanced QoS features that ensure complete isolation among the multiple hosts connected to the NIC, and by achieving superior fairness among the hosts.

Solutions

- > Cloud-native, Web 2.0, hyperscale
- > Enterprise data centers
- > Cybersecurity
- > Big data analytics
- > Scale-out compute and storage infrastructure
- > Telco and network function virtualization (NFV)
- > Cloud storage
- > Machine learning and AI
- > Media and entertainment

FEATURES*

Network Interface

- > 2 x 1/10/25/40/50/100 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 1K virtual functions per port
 - > 8 physical functions
- > Support for tunneling
 - > Encap/decap of VXLAN, NVGRE, Geneve, and more
 - > Stateless offloads for overlay tunnels

NVIDIA 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 256 /512-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 packets
- > Transmit side scaling (TSS)
- > VLAN and MPLS tag insertion/stripping
- > Receive flow steering

Storage Offloads

- > Block-level encryption: XTS-AES 256/512-bit key
- > NVMe over Fabric offloads for target machine
- > T10 DIF - signature handover operation at wire speed, for ingress and egress traffic
- > Storage protocols: SRP, iSER, NFS RDMA, SMB Direct, NVMe-oF

Advanced Timing & Synchronization

- > Advanced PTP
 - > IEEE 1588v2 (any profile)
 - > PTP hardware clock (PHC) (UTC format)
 - > 20nsec 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)

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, NCSI over RBT in OCP 2.0/3.0 cards
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP026
- > I²C 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 Gen3 and Gen4

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

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 ¹ , With Secure Boot
2 x 25GbE	SFP28	1/10/25	Gen4 x8	MCX621102AN-ADAT	MCX621102AC-ADAT
			Gen4 x16	MCX623102AN-ADAT	MCX623102AC-ADAT
2 x 50GbE	SFP56	1/10/25/50	Gen4 x16	MCX623102AN-GDAT	MCX623102AC-GDAT
1 x 100GbE	QSFP56	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623105AN-CDAT	Contact NVIDIA
2 x 100GbE	QSFP56	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623106AN-CDAT	MCX623106AC-CDAT
			Gen4 2x8 in a row, NVIDIA Socket Direct	Contact NVIDIA	Contact NVIDIA
	QSFP56 + PPS In/Out SMAs	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623106PN-CDAT	MCX623106PC-CDAT
	DSFP	1/10/25/40/50 ² /100	Gen4 x16	Contact NVIDIA	Contact NVIDIA
1 x 200GbE	QSFP56	10/25/40/50 ² /100 ³ /200	Gen4 x16	MCX623105AN-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.



QCP 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 ¹ , With Secure Boot
2 x 25GbE	SFP28	1/10/25	Gen4 x16	MCX623432AN-ADAB	MCX623432AC-ADAB
2 x 50GbE	SFP56	1/10/25/50	Gen4 x16	MCX623432AN-GDAB	MCX623432AC-GDAB
1 x 100GbE	QSFP56	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623435AN-CDAB	MCX623435AC-CDAB
			Gen4 x16 Multi-Host	MCX623435MN-CDAB	Contact NVIDIA
	DSFP	1/10/25/50/100	Gen4 x16 Multi-Host	Contact NVIDIA	MCX623439MC-CDAB
2 x 100GbE	QSFP56	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623436AN-CDAB	MCX623436AC-CDAB
			Gen4 x16 Multi-Host	MCX623436MN-CDAB	Contact NVIDIA
	DSFP	1/10/25/50/100	Gen4 x16 Multi-Host	Contact NVIDIA	Contact NVIDIA
1 x 200GbE	QSFP56	1/10/25/40/50 ² /100 ³ /200	Gen4 x16	MCX623435AN-VDAB	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 NVIDIA for OCP OPNs with NVIDIA Multi-Host support.

The above QCP3.0 OPNs come with thumbscrew (pull tab) brackets; contact NVIDIA for additional bracket types, such as Internal Lock or Ejector latch.



QCP 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 ¹ , With Secure Boot
1 x 100GbE	QSFP56	1/10/25/40/50 ² /100 ³	Gen4 x16	MCX623405AN-CDAN	MCX623405AC-CDAN
1 x 200GbE	QSFP56	1/10/25/40/50 ² /100 ³ /200	Gen4 x16	MCX623435AN-VDAB	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 NVIDIA for OCP OPNs with NVIDIA Multi-Host or NVIDIA Socket Direct support.

ConnectX-6 Dx PCIe Gen4 cards are backward compatible.



*For illustration only. Actual products may vary.

Learn more at www.nvidia.com/en-us/networking/ethernet/connectx-6-dx

© 2021 Mellanox Technologies. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, ConnectX, NVIDIA Multi-Host, NVIDIA Socket Direct, NVIDIA 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. MAY21/60259PB-R5

