

NVIDIA CONNECTX-6 DX CISCO 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, big data, machine learning, and storage. It provides two ports of 100Gb/s Ethernet connectivity and the highest ROI of any SmartNIC.

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 and AES-XTS block-level, data-at-rest encryption and decryption offloads
- > Probes and denial-of-service (DoS) attack protection—**A hardware-based L4 firewall** is achieved by offloading stateful connection tracking through NVIDIA ASAP² - Accelerated Switch and Packet Processing®
- > NIC security—hardware root-of-trust (RoT) secure boot and firmware updates 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

- > 100Gb/s Ethernet, 50Gb/s pulse-amplitude modulation 4-level (PAM4), and 25Gbp/s / 10Gb/s non-return to zero (NRZ) SerDes
- > PCIe Gen3 and Gen4 x16 host interface
- > Crypto and non-crypto versions

Key Features

- > ASIC capable of up to 200Gb/s of full duplex bandwidth
- > Message rate of up to 215Mpps
- > Sub 0.8usec latency
- > Programmable pipeline for new network flows
- > NVIDIA Multi-Host™ with advanced QoS
- > Auto negotiation with 40G switches
- > Auto selection of PAM4 and NRZ according to selected speed
- > 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 updates
- > Connection tracking offload
- > Advanced RoCE capabilities

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 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 and time-based, software-defined networking (SDN) accelerations. 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.

EFFICIENT STORAGE SOLUTIONS

With its Non-Volatile Memory Express over Fabrics (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.

ORDERING INFORMATION

| Cisco Product ID | NVIDIA Part Number | Description | Supported Ethernet Speeds [GbE] | Crypto and/or Secure Boot Support | Qualified Cisco Servers |
|---------------------------------------|--------------------|---|---------------------------------|---|-------------------------------|
| UCSC-P-M6CD100GF UCSC-P-M6CD100GF= | MCX62310 6AC-CDAT | ConnectX-6 Dx EN adapter card, 100GbE, dual-port QSFP56, PCIe 4.0 x16, crypto and secure boot, tall bracket | 1/10/25/40/50/100 ² | Crypto supported; secure boot supported | Cisco UCS C220 M6 rack server |
| UCSC-P-M6DD100GF UCSC-P-M6DD100GF= | MCX62310 6AS-CDAT | ConnectX-6 Dx EN adapter card, 100GbE, dual-port QSFP56, PCIe 4.0 x16, secure boot, no crypto, tall bracket | 1/10/25/40/50/100 ² | Secure boot supported; no crypto | Cisco UCS C240 M6 rack server |

¹ Use crypto-enabled cards to utilize IPsec/TLS/AES-XTS encryption/decryption hardware offloads.

² 100G can be supported as either 4x 25G NRZ or 2x 50G PAM4 when using QSFP56.

Cards come assembled with a tall bracket PN MEC018771, if needed the short bracket PN MEC016919 is available in the box as well.

"=" signifies a Cisco spare part is orderable separately.

Server qualification is presented as of the date of publication. For latest server qualification information, please visit the [Cisco Hardware Compatibility List \(HCL\)](#).

- > 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

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; using QSFP to SFP (QSA) modules will allow use of lower speeds, see **MAM1Q00A-QSA28**.
- > Automatic link configuration validated with NVIDIA 1G, 10G, and 25G switches
- > Manual link configuration may be required with non-NVIDIA switches

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-XTS 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 packet
- > Transmit side scaling (TSS)
- > VLAN and MPLS tag insertion/stripping
- > Receive flow steering

Storage

- > Block-level encryption offload: 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, NVMe over TCP

Advanced Timing and 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

Operating Systems/Distributions

- > RHEL, SLES, Ubuntu, and other major Linux distributions
- > Windows
- > FreeBSD
- > VMware
- > CentOS

Management and Control

- > MCTP over SMBus and MCTP over PCIe - Baseboard Management Controller interface
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP0267
- > 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 802.3cd, 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 "FastWake" mode)
- > IEEE 802.3ap based autonegotiation 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.

OPTICS AND CABLES

NVIDIA recommends using NVIDIA cables and modules. **Additional information on tested modules:**
 Select firmware version > Select NVIDIA OPN > Select PSID > Select “Release Notes” under Download/Documentation.

CISCO-BRANDED SUPPORTED CABLES AND MODULES

| QSFP Product ID | Product Description |
|------------------|---|
| QSFP-100G-AOC10M | 100GBASE QSFP active optical cable, 10m |
| QSFP-100G-AOC7M | 100GBASE QSFP active optical cable, 7m |
| QSFP-100G-AOC5M | 100GBASE QSFP active optical cable, 5m |
| QSFP-100G-CU5M | 100GBASE-CR4 passive copper cable, 5m |
| QSFP-100G-CU3M | 100GBASE CR4 passive copper cable, 3m |
| QSFP-100G-LR4-S | 100GBASE-LR4 QSFP transceiver, LC, 10km over SMF |
| QSFP-100G-SR4-S | 100GBASE SR4 QSFP transceiver, MPO, 100m over OM4 MMF |
| QSFP-40/100-SRBD | 100G and 40GBASE SR-BiDi QSFP transceiver, LC, 100m OM4 MMF |
| QSFP-100G-DR-S | 100GBASE DR QSFP transceiver, 500m over SMF |
| QSFP-100G-LR4-S | 100GBASE LR4 QSFP transceiver, LC, 10km over SMF |

For latest updates, check the UCS server technical specs; also consult the [Cisco Transceiver Module Group \(TMG\) Compatibility Matrix](#). Break-out and split cables are not supported.

ENVIRONMENTAL SPECIFICATIONS

| Temperature | |
|--|---|
| Operating: 0°C to 55°C (32°F to 131°F) | Storage: -40°C to 70°C (-40°F to 158°F) |

COMPATIBILITY-TESTED SWITCHES

| Company and Model | PAM4 or NRZ Testing |
|-------------------|---------------------|
| Cisco C9236C | NRZ |
| Cisco C3232C | NRZ |
| NVIDIA SN3700 | PAM4 and NRZ |
| NVIDIA SN2740 | NRZ |
| NVIDIA SN2410 | NRZ |
| NVIDIA SX1036 | NRZ |

The cards will automatically select PAM4 or NRZ mode based on the link configuration with the switch.

AIRFLOW REQUIREMENTS

| Cable | Hot Aisle—Heatsink to Port |
|-------------------|----------------------------|
| Passive cable | 550LFM @ 55°C |
| Active 2.5W cable | 700LFM @ 55°C |
| Active 3.5W cable | 1100LFM @ 55°C |

Airflow is measured in a wind tunnel. Contact NVIDIA for airflow numbers with other active module power levels.

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

© 2021 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, ASAP² - Accelerated Switch and Packet Processing, ConnectX, GPUDIRECT, and NVIDIA Multi-Host are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.
 MAY21 60259C-R5

