

# NVIDIA CONNECTX-6 DX

## Ethernet SmartNIC

### Advanced Networking and Security for the Most Demanding Cloud and Data Center Workloads

NVIDIA® ConnectX®-6 Dx is a highly secure and advanced smart network interface card (SmartNIC) that accelerates mission-critical cloud and data center applications, including security, virtualization, SDN/NFV, big data, machine learning, and storage. ConnectX-6 Dx provides up to two ports of 100Gb/s or a single port of 200Gb/s Ethernet connectivity and is powered by 50Gb/s (PAM4) or 25/10 Gb/s (NRZ) SerDes technology.

ConnectX-6 Dx features virtual switch (vSwitch) and virtual router (vRouter) hardware accelerations delivering orders-of-magnitude higher performance than software-based solutions. ConnectX-6 Dx supports a choice of single-root I/O virtualization (SR-IOV) and VirtIO in hardware, enabling customers to best address their application needs. By offloading cloud networking workloads, ConnectX-6 Dx frees up CPU cores for business applications while reducing total cost-of-ownership.

In an era where data privacy is key, ConnectX-6 Dx provides built-in inline encryption/decryption, stateful packet filtering, and other capabilities, bringing advanced security down to every node with unprecedented performance and scalability.


Built on the solid foundation of NVIDIA's ConnectX line of SmartNICs, ConnectX-6 Dx offers best-in-class RDMA over Converged Ethernet (RoCE) capabilities, enabling scalable, resilient, and easy-to-deploy RoCE solutions. For data storage, ConnectX-6 Dx optimizes a suite of storage accelerations, bringing NVMe-oF target and initiator offloads.

### PRODUCT SPECIFICATIONS

Maximum total bandwidth	<b>200Gb/s</b>
Supported Ethernet speeds	<b>10/25/40/50/100/200GbE</b>
Number of network ports	<b>1/2</b>
Network interface technologies	<b>NRZ/PAM4</b>
Host interface	<b>PCIe Gen4.0 x16, with NVIDIA Multi-Host™ technology</b>
DPDK message rate	<b>Up to 215Mpps</b>
Platform security	<b>Hardware root-of-trust and secure firmware update</b>
Form factors	<b>PCIe HHHL, OCP2, OCP3.0 SFF</b>
Network interfaces	<b>SFP+, QSFP+, DSFP</b>

### 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



PCIe x16 HHHL Card



OCP 3.0 Small Form Factor



OCP 2.0 Form Factor

\*For illustration only. Actual products may vary.

## Features(\*)

### Network Interface

- > Dual ports of 10/25/40/50/100 GbE, or a single port of 200GbE

### Host Interface

- > 16 lanes of PCIe Gen4, compatible with PCIe Gen2/Gen3
- > Integrated PCI switch
- > NVIDIA Multi-Host and NVIDIA Socket Direct™

### Virtualization/Cloud Native

- > 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 ASAP<sup>2</sup> Accelerated Switching & Packet Processing

- > SDN acceleration for:
  - > Bare metal
  - > Virtualization
  - > Containers
- > Full hardware offload for OVS data plane
- > Flow update through RTE\_Flow or TC\_Flower
- > Flex-parser: user-defined classification
- > Hardware offload for:
  - > Connection tracking (Layer 4 firewall)
  - > NAT
  - > Header rewrite
  - > Mirroring
  - > Sampling
  - > Flow aging
  - > Hierarchical QoS
  - > Flow-based statistics

### Cybersecurity

- > Inline hardware IPsec encryption and decryption
  - > AES-GCM 128/256-bit key
  - > RoCE over IPsec
- > 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 Offloads

- > Block-level encryption: XTS-AES 256/512-bit key
- > NVMe over Fabrics 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 and Synchronization

- > Advanced PTP
  - > IEEE 1588v2 (any profile)
  - > PTP hardware clock (PHC) (UTC format)
  - > Nanosecond-level 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)
- > Dedicated precision timing card option

### RDMA over Converged Ethernet (RoCE)

- > RoCE v1/v2
- > Zero-touch RoCE: no ECN, no PFC
- > RoCE over overlay networks
- > Selective repeat
- > Programmable congestion control interface
- > GPUDirect®

### Management and Control

- > NC-SI, MCTP over SMBus and MCTP over PCIe—Baseboard Management Controller interface, NCSI over RBT in Open Compute Project (OCP) 2.0/3.0 cards
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP0267
- > I<sup>2</sup>C interface for device control and configuration

### Remote Boot

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

## Ordering Information

For NVIDIA ordering information, please contact your NVIDIA sales representative or visit the online ConnectX-6 Dx user manuals: **PCIe HHHL form factor**, **OCP 3.0 form factor** and **OCP 2.0 form factor**.

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

[Learn more](#)

Learn More at [NVIDIA.com/en-us/networking/ethernet/connectx-6-dx](https://www.nvidia.com/en-us/networking/ethernet/connectx-6-dx)

© 2021 NVIDIA Corporation & Affiliates. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, ConnectX, GPUDirect, Multi-Host, Socket Direct, and ASAP<sup>2</sup> - Accelerated Switch & Packet Processing are trademarks and/or registered trademarks of NVIDIA Corporation and its affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. All other trademarks are property of their respective owners. OCT21/60259PB-R6.

