

ConnectX[®] 5

ConnectX[®] -5 EN Card

Up to 100Gb/s Ethernet Adapter Card

Intelligent RDMA-enabled, single and dual-port SmartNIC network adapter with advanced application offload capabilities for Web2.0, Cloud, Storage, and Telco platforms

ConnectX-5 EN adapter cards provide high performance and flexible solutions, with up to two ports of 100GbE connectivity, 750ns latency, up to 200 million messages per second (Mpps), and a record-setting 150 Mpps when running an open source Data Path Development Kit (DPDK) PCIe. For storage workloads, ConnectX-5 delivers a range of innovative accelerations, such as Signature Handover (T10-DIF) in hardware, an embedded PCIe Switch, and NVMe over Fabric Targets offloads. ConnectX-5 NICs also bring advanced OvS Offloads to telecommunications and cloud data centers to drive extremely high packet rate and throughput with reduced CPU resource consumption, thus boosting data center infrastructure efficiency.

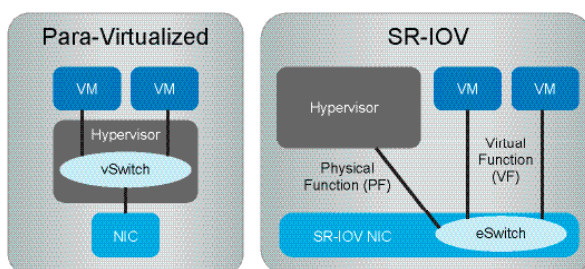
ConnectX-5 network adapters are available for PCIe Gen3 and Gen4 servers (ConnectX-5 Ex) and provide support for 1, 10, 25, 40, 50 and 100GbE speeds in standup PCIe card, OCP 2.0, 3.0 and OEM customized form factors. ConnectX-5 adapter cards also offer advanced Mellanox Multi-Host[®] and Mellanox Socket Direct[®] technologies.

Cloud and Web2.0 Environments

ConnectX-5 NICs enable data center administrators to benefit from better server utilization and reduced costs, power usage, and cable complexity, allowing for more virtual appliances, virtual machines (VMs) and tenants to co-exist on the same hardware.

Supported vSwitch/vRouter offload functions include:

- Overlay Networks (e.g., VXLAN, NVGRE, MPLS, GENEVE, and NSH) header encapsulation & de-encapsulation.
- Stateless offloads of inner packets and packet headers' re-write, enabling NAT functionality and more.
- Flexible and programmable parser and match-action tables, which enable hardware offloads for future protocols.
- SR-IOV technology, providing dedicated adapter resources, guaranteed isolation and protection for virtual machines (VMs) within the server.
- Network Function Virtualization (NFV), enabling a VM to be used as a virtual appliance. The full data-path operation offloads, hairpin hardware capability and service chaining enables data to be handled by the Virtual Appliance, with minimum CPU utilization.



FEATURES

- Tag matching and rendezvous offloads
- Adaptive routing on reliable transport
- Burst buffer offloads for background checkpointing
- NVMe over Fabric (NVMe-oF) offloads
- Back-end switch elimination by host chaining
- Embedded PCIe switch
- Enhanced vSwitch/vRouter offloads
- Flexible pipeline
- RoCE for overlay networks
- PCIe Gen 4 support
- RoHS-compliant
- ODCC-compatible

BENEFITS

- Up to 100Gb/s connectivity per port
- Industry-leading throughput, low latency, low CPU utilization and high message rate
- Innovative rack design for storage and Machine Learning based on Host Chaining technology
- Smart interconnect for x86, Power, Arm, and GPU-based compute and storage platforms
- Advanced storage capabilities including NVMe over Fabric offloads
- Intelligent network adapter supporting flexible pipeline programmability
- Cutting-edge performance in virtualized networks including Network Function Virtualization (NFV)
- Enabler for efficient service chaining capabilities
- Efficient I/O consolidation, lowering data center costs and complexity

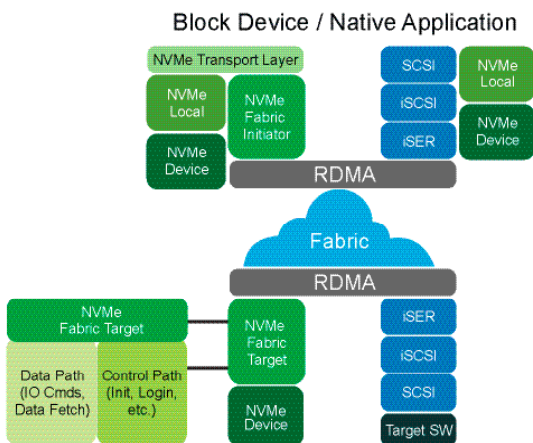
Cloud and Web2.0 customers developing platforms on Software Defined Network (SDN) environments are leveraging their servers' Operating System Virtual-Switching capabilities to achieve maximum flexibility. Open V-Switch (OvS) is an example of a virtual switch that allows Virtual Machines to communicate with each other and with the outside world. Traditionally residing in the hypervisor where switching is based on twelve-tuple matching onflows, the virtual switch, or virtual router software-based solution, is CPU-intensive. This can negatively affect system performance and prevent the full utilization of available bandwidth.

Mellanox ASAP² - Accelerated Switching and Packet Processing[®] technology enables offloading the vSwitch/vRouter by handling the data plane in the NIC hardware, without modifying the control plane. This results in significantly higher vSwitch/vRouter performance without the associated CPU load.

Additionally, intelligent ConnectX-5's flexible pipeline capabilities, including flexible parser and flexible match-action tables, are programmable, enabling hardware offloads for future protocols.


Storage Environments

NVMe storage devices are gaining popularity by offering very fast storage access. The evolving NVMe over Fabric (NVMe-oF) protocol leverages the RDMA connectivity for remote access. ConnectX-5 offers further enhancements by providing NVMe-oF target offloads, enabling very efficient NVMe storage access with no CPU intervention, and thus improving performance and reducing latency.



The embedded PCIe switch enables customers to build standalone storage or Machine Learning appliances. As with the earlier generations of ConnectX adapters, standard block and file access protocols leverage RoCE for high-performance storage access. A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks.

ConnectX-5 enables an innovative storage rack design, Host Chaining, which enables different servers to interconnect without involving the Top of the Rack (ToR) switch. Leveraging Host Chaining, ConnectX-5 lowers the data center's Total Cost of Ownership (TCO) by reducing CAPEX (cables, NICs, and switch port expenses). Also OPEX is reduced by cutting down on switch port management and overall power usage.



Telecommunications

Telecommunications service providers are moving towards disaggregation, server virtualization, and orchestration as key tenets to modernize their networks. Likewise, they're also moving towards Network Function Virtualization (NFV), which enables the rapid deployment of new network services. With this move, proprietary dedicated hardware and software, which tend to be static and difficult to scale, are being replaced with virtual machines running on commercial off-the-shelf (COTS) servers.

For Telecom service providers, choosing the right networking hardware is critical to achieving a cloud-native NFV solution that is agile, reliable, fast and efficient. Telco service providers typically leverage virtualization and cloud technologies to better achieve agile service delivery and efficient scalability; these technologies require an advanced network infrastructure to support higher rates of packet processing. However, the resultant east-west traffic causes numerous interrupts as I/O traverses from kernel to user space, eats up CPU cycles and decreases packet performance. Particularly sensitive to delays are voice and video applications which often require less than 100ms of latency.

Mellanox's ConnectX-5 adapters drive extremely high packet rates, increased throughput and drive higher network efficiency through the following technologies; Open vSwitch Offloads (OvS), OvS over DPDK or ASAP², Network Overlay Virtualization, SR-IOV, and RDMA. This allows for secure data deliver through higher-performance offloads, reducing CPU resource utilization, and boosting data center infrastructure efficiency. The result is a much more responsive and agile network capable of rapidly deploying network services.

Host Management

Mellanox host management and control capabilities include NC-SI over MCTP over SMBus, and MCTP over PCIe - Baseboard Management Controller (BMC) interface, as well as PLDM for Monitor and Control DSP0248 and PLDM for Firmware Update DSP0267.

Compatibility

PCI Express Interface

- PCIe Gen 4
- PCIe Gen 3.0, 1.1 and 2.0 compatible
- 2.5, 5.0, 8, 16GT/s link rate
- Auto-negotiates to x16, x8, x4, x2, or x1 lanes
- PCIe Atomic
- TLP (Transaction Layer Packet) Processing Hints (TPH)
- Embedded PCIe Switch: Up to 8 bifurcations

- PCIe switch Downstream Port Containment (DPC) enablement for PCIe hot-plug
- Access Control Service (ACS) for peer-to-peer secure communication
- Advance Error Reporting (AER)
- Process Address Space ID (PASID) Address Translation Services (ATS)
- IBM CAPI v2 support (Coherent Accelerator Processor Interface)
- Support for MSI/MSI-X mechanisms

Operating Systems/Distributions*

- RHEL/CentOS
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF-2)

Connectivity

- Interoperability with Ethernet switches (up to 100GbE)
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

Features

Ethernet

- 100GbE / 50GbE / 40GbE / 25GbE / 10GbE / 1GbE
- IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet
- IEEE 802.3by, Ethernet Consortium 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 (fast wake)
- 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 802.1Qau (QCN) – Congestion Notification
- IEEE 802.1Qaz (ETS)
- IEEE 802.1Qbb (PFC)
- IEEE 802.1Qbg
- IEEE 1588v2
- Jumbo frame support (9.6KB)

Enhanced Features

- Hardware-based reliable transport
- Collective operations offloads
- Vector collective operations offloads
- PeerDirect™ RDMA (aka GPUDirect®) communication acceleration
- 64/66 encoding
- Extended Reliable Connected transport (XRC)

- Dynamically Connected Transport (DCT)
- Enhanced Atomic operations
- Advanced memory mapping support, allowing user mode registration and remapping of memory (UMR)
- On demand paging (ODP)
- MPI Tag Matching
- Rendezvous protocol offload
- Out-of-order RDMA supporting Adaptive Routing
- Burst buffer offload
- In-Network Memory registration-free RDMA memory access

CPU Offloads

- RDMA over Converged Ethernet (RoCE)
- TCP/UDP/IP stateless offload
- LSO, LRO, checksum offload
- RSS (also on encapsulated packet), TSS, HDS, VLAN and MPLS tag insertion/stripping, Receive flow steering
- Data Plane Development Kit (DPDK) for kernel bypass applications
- Open vSwitch (OVS) offload using ASAP²
 - Flexible match-action flow tables
 - Tunneling encapsulation/de-capsulation
- Intelligent interrupt coalescence
- Header rewrite supporting hardware offload of NAT router

Storage Offloads

- 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

Overlay Networks

- RoCE over Overlay Networks
- Stateless offloads for overlay network tunneling protocols
- Hardware offload of encapsulation and decapsulation of VXLAN, NVGRE, and GENEVE overlay networks

Hardware-Based I/O Virtualization

- Single Root IOV
- Address translation and protection
- VMware NetQueue support
- SR-IOV: Up to 1K Virtual Functions
- SR-IOV: Up to 16 Physical Functions per host
- Virtualization hierarchies (e.g., NPAR when enabled)
 - Virtualizing Physical Functions on a physical port
 - SR-IOV on every Physical Function
- Configurable and user-programmable QoS
- Guaranteed QoS for VMs

HPC Software Libraries

- Open MPI, IBM PE, OSU MPI (MVAPICH/2), Intel MPI
- Platform MPI, UPC, Open SHMEM

Management and Control

- NC-SI over MCTP over SMBus and NC-SI over MCTP over PCIe - Baseboard Management Controller interface
- PLDM for Monitor and Control DSP0248
- PLDM for Firmware Update DSP0267
- SDN management interface for managing the eSwitch
- 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
- Unified Extensible Firmware Interface (UEFI)
- Pre-execution Environment (PXE)

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

Table 1 - Part Numbers and Descriptions

OPN	Ethernet Supported Speeds (GbE)	Network Ports	Cage(s)	PCI Express Configuration	Form Factor	Additional Features
MCX511F-ACAT	25,10,1	1	SFP28	PCIe 3.0 x16	PCIE Standup	
MCX512A-ACAT	25,10,1	2	SFP28	PCIe 3.0 x8	PCIE Standup	
MCX512A-ACUT	25,10,1	2	SFP28	PCIe 3.0 x8	PCIE Standup	UEFI Enabled (x86/ARM),
MCX512A-ADAT	25,10,1	2	SFP28	PCIe3.0/4.0 x8	PCIE Standup	Improved Latency
MCX512F-ACAT	25,10,1	2	SFP28	PCIe 3.0 x16	PCIE Standup	
MCX512F-ACHT	25,10,1	2	SFP28	PCIe 3.0 x16	PCIE Standup	Enhanced Host Management
MCX516A-BDAT	40,25,10,1	2	QSFP28	PCIe 4.0 x16	PCIE Standup	
MCX515A-GCAT	50,40,25,10,1	1	QSFP28	PCIe 3.0 x16	PCIE Standup	
MCX516A-GCAT	50,40,25,10,1	2	QSFP28	PCIe 3.0 x16	PCIE Standup	
MCX515A-CCAT	100,50,40,25,10,1	1	QSFP28	PCIe 3.0 x16	PCIE Standup	
MCX516A-CCAT	100,50,40,25,10,1	2	QSFP28	PCIe 3.0 x16	PCIE Standup	
MCX516A-CCHT	100,50,40,25,10,1	2	QSFP28	PCIe 3.0 x16	PCIE Standup	Enhanced Host Management
MCX516A-CDAT	100,50,40,25,10,1	2	QSFP28	PCIe 4.0 x16	PCIE Standup	

NOTES:

- (1) All tall-bracket adapters are shipped with the tall bracket mounted and a short bracket as an accessory.
- (2) Dimensions without bracket: 14.2cm x 6.9cm (Low Profile)



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com