

ConnectX[®] InfiniBand Smart Adapter Cards for OCP

High Performance Smart Adapter Cards supporting up to 200Gb/s InfiniBand & Ethernet in Open Compute Project Spec 2.0 & 3.0 Form Factors

Mellanox® Virtual Protocol Interconnect® (VPI) smart adapter cards in OCP 2.0 & 3.0 form factors combine high-performance leading features with best-in-class efficiency, enabling the highest data center performance.

World-Class Performance and Scale

Mellanox smart adapter cards deliver industry-leading connectivity for performance-driven server and storage applications. ConnectX adapter cards offer high bandwidth coupled with ultra-low latency and in-network computing to enable faster access and real-time responses.

Mellanox offers a variety of OCP Spec 2.0 and OCP Spec 3.0 compliant adapter cards, providing best-in-class performance and efficient computing through advanced acceleration and offload capabilities. These advanced capabilities free up valuable CPU cores for other tasks, while increasing data center performance, scalability and efficiency, include:

- IBTA Remote Data Memory Access (RDMA)
- NVMe-over-Fabrics (NVMe-oF) offloads
- In-network computing and MPI operation accelerations
- Single Root IO Virtualization (SR-IOV)
- GPUDirect® communication acceleration
- Mellanox Multi-Host® for connecting multiple compute or storage hosts to a single interconnect adapter card
- Enhanced security solutions

Complete End-to-End InfiniBand Networking

ConnectX OCP smart adapter cards are part of Mellanox's InfiniBand end-to-end portfolio for data centers which also includes switches, application acceleration packages, and cabling to deliver the best price-performance value proposition for network and storage solutions. With Mellanox, IT managers can be assured of the highest performance, reliability, most efficient network fabric at the lowest cost, and the best return on investment.

In addition, Mellanox Unified Fabric Manager (UFM®) management software greatly simplifies network provisioning, monitoring and diagnostics, providing the agility and efficiency for scalability and future growth. Featuring an intuitive graphical user interface, UFM provides in-depth visibility and control.

Open Compute Project

The OCP NIC 3.0 specification extends the capabilities of OCP NIC 2.0 design specification. OCP 3.0 defines a different form factor and connector style than OCP 2.0. The OCP 3.0 specification defines two basic card sizes: Small Form Factor (SFF) and Large Form Factor (LFF). Mellanox OCP NICs are currently supported in a SFF.*

* Future designs may utilize LFF to allow for additional PCIe lanes and/or network ports,

OCP 3.0 provides additional board real estate, thermal capacity, electrical interfaces, network interfaces, host flagnration and management. OCP 3.0 also introduces a new mating technique that simplifies FRU installation and removal, and reduces overall downtime.

The table below shows key comparisons between the OCP Specs 2.0 and 3.0.

Feature	OCP Spec 2.0	OCP Spec 3.0
Card Dimensions	Non-rectangular (8000mm ²)	SFF: 76x115mm (8740mm ²)
PCIe Lanes	Up to x16	SFF: Up to x16
Maximum Power Capability	Up to 67.2W for PCIe x8 card; Up to 86.4W for PCIe x16 card	SFF: Up to 80W
Baseband Connector Type	Mezzanine (B2B)	Edge (0.6mm pitch)
Network Interfaces	Up to 2 SFP side-by-side or 2 QSFP belly-to-belly	Up to two QSFP in SFF, side-by-side
Expansion Direction	N/A	Side
Installation in Chassis	Parallel to front or rear panel	Perpendicular to front/rear panel
Hot Swap	No	Yes (pending server support)
Multi-Host	Up to 4 hosts	Up to 4 hosts in SFF or 8 hosts in LFF
Host Management Interfaces	RBT, SMBus	RBT, SMBus, PCIe
Host Management Protocols	Not standard	DSP0267, DSP0248

For more details, please refer to the Open Compute Project (OCP) Specifications.

ConnectX OCP Smart Adapter Card Benefits

- Open Data Center Committee (ODCC) compatible
- Compliant with OCP 2.0 & 3.0 NIC specifications
- All platforms: x86, Power, Arm, compute and storage
- Industry-leading performance
- RDMA and TCP/IP for I/O consolidation
- SR-IOV technology: VM protection and QoS
- Increased Virtual Machine (VM) count per server

Target Applications

- High Performance Computing (HPC)
- Machine Learning, Artificial Intelligence (AI)
- Compute and storage platforms
- Data center virtualization
- Clustered databases and high-throughput data warehousing
- Latency-sensitive financial analysis and high frequency trading



InfiniBand OCP Smart Adapter Cards

Specs, Feature Support & Ordering Part Numbers



OCP Spec & Form Factor	InfiniBand Network Speeds	Ethernet Network Speeds	Network Interface(s)	PCIe Host Interface	Product Family	Mellanox Multi-Host	OPN
OCP 2.0 Type 2	Up to 100Gb/s EDR, FDR, QDR, SDR	100, 50, 40, 25, 10 GbE	1x QSFP28	Gen 3.0 x16	ConnectX-5 VPI	--	MCX545A-ECAN
			1x QSFP28	Gen 3.0 x16		✓	MCX545M-ECAN
			2x QSFP28 Belly-to-Belly	Gen 4.0 x16		--	MCX546A-EDAN
OCP 2.0 Type 1	Up to 100Gb/s EDR, FDR, QDR, SDR	100, 50, 40, 25, 10 GbE	1x QSFP28	Gen 4.0 x16		--	MCX545B-ECAN
OCP 3.0 SFF	Up to 200Gb/s HDR, EDR, FDR, QDR, SDR	200, 100, 50, 40, 25, 10 GbE	1x QSFP56	Gen 3.0/4.0 x16	ConnectX-6 VPI	--	MCX653435A-HDAI
			1x QSFP56	Gen 3.0/4.0 x16		--	MCX653435A-HDAE
			2x QSFP56	Gen 3.0/4.0 x16		--	MCX653436A-HDAI
	Up to 100Gb/s HDR100, EDR, FDR, QDR, SDR	100, 50, 40, 25, 10 GbE	1x QSFP56	Gen 3.0/4.0 x16		--	MCX653435A-EDAI

Notes:

- Additional information is available in product briefs.
- This brochure describes hardware features and capabilities. Please refer to the driver and firmware release notes on www.mellanox.com for feature availability.
- OPN2.0 Product Notes:
 - Shipped without a bracket.
 - For more details, please refer to the Open Compute Project 2.0 Specifications.
- OCP3.0 Product Notes:
 - The last digit of OPN-suffix displays the default bracket option: I = Internal Lock; E = Ejector Latch. For other bracket types, contact Mellanox.
 - Dimensions without brackets are 167.65mm x 68.90mm. All tall-bracket adapters are shipped with the tall bracket mounted and a short bracket as an accessory.



Enabling High Performance Computing (HPC) and Artificial Intelligence (AI) Applications

Mellanox InfiniBand/VPI adapters are the perfect solution for the evolving data-centric paradigm. Technologies within this model include the innovative In-Network Computing offloads that transform the data center interconnect into a “distributed CPU,” and “distributed memory,” overcoming performance bottlenecks and enabling faster and more scalable data analysis.

Mellanox’s advanced In-Network Computing accelerations and RDMA offload capabilities optimize the performance of a wide variety of HPC and machine learning systems in bioscience, media, automotive design, CFD and manufacturing, weather research, oil and gas, and other markets.

As a core In-Networking Computing technology, Mellanox Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™] optimizes MPI operations and GPU communications performance, decreasing the data load on the network and dramatically reducing operation times, while freeing up CPU and/or GPU resources for other tasks.

Virtual Protocol Interconnect[®] (VPI)

Mellanox’s VPI technology enables any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. Each port can operate on InfiniBand or Ethernet fabrics, and supports IP over InfiniBand (IPoB) and RDMA over Converged Ethernet (RoCE). VPI simplifies I/O system design and makes it easier for IT managers to deploy infrastructure that meets the challenges of a dynamic data center.

I/O Virtualization

Mellanox adapter cards provide comprehensive support for virtualized data centers with Single-Root I/O Virtualization (SR-IOV) allowing dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization on InfiniBand gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

Mellanox Multi-Host Solution

Mellanox Multi-Host[®] technology provides high flexibility and major savings in building next generation, scalable, high-performance data centers. Mellanox Multi-Host connects multiple compute or storage hosts into a single interconnect adapter, separating the adapter’s PCIe interface into multiple and independent PCIe interfaces with no performance degradation. The technology enables designing and building new scale-out heterogeneous compute and storage racks with direct connectivity between compute elements, storage elements and the network, better power and performance management, while achieving maximum data processing and data transfer at minimum capital and operational expenses.

Accelerated Storage

A consolidated compute and storage network provides significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols leveraging InfiniBand RDMA result in high-performance storage access. Mellanox adapter cards support NVMe over Fabrics as well as other technologies including SRP, iSER, NFS RDMA, SMB Direct, SCSI and iSCSI. ConnectX adapters also offer a flexible Signature Handover mechanism based on the advanced T-10/DIF implementation.

ConnectX[®]-6 Adapter Card

ConnectX-6 is the world's first 200Gb/s HDR InfiniBand and Ethernet network adapter card, offering world-leading performance, smart offloads and In-Network Computing. ConnectX-6 with VPI provides two ports of 200Gb/s supporting HDR, HDR100, EDR, FDR, QDR and SDR InfiniBand speeds, as well as 200, 100, 50, 40, 25 and 10 Gb/s Ethernet speeds. ConnectX-6 achieves super-low latency and up to 215 million messages/second.

In addition to all the features included in earlier models of ConnectX, ConnectX-6 adapter cards may offer Mellanox Multi-Host support for up to 8 hosts and block-level encryption as a crucial innovation to network security, altogether delivering the highest performance, most secure and extremely flexible solution for today's demanding applications and markets.

ConnectX[®]-5 Adapter Card

Intelligent ConnectX-5 adapter cards support co-design and in-network computing, while introducing acceleration engines for maximizing HPC, data analytics and storage platforms. They support two ports of up to EDR 100Gb/s InfiniBand and 100Gb/s Ethernet connectivity, super low latency, a very high message rate, plus PCIe switch and NVMe over Fabric offloads. In addition, ConnectX-5 is capable of supporting Message Passing Interface (MPI) offloads, such as MPI Tag Matching, and enables switches' adaptive routing capabilities.

(For specific adapter card models refer to the table on page 4.)

Broad Software Support

All Mellanox adapter cards are supported by a full suite of drivers for Linux major distributions, Microsoft[®] Windows[®], VMware vSphere[®] and FreeBSD[®].

Drivers are also available in-box in Linux main distributions, Windows and VMware.



Multiple Form Factors

Mellanox smart adapter cards are available in a variety of form factors to meet specific data center needs, including:


- OCP Spec 2.0 Type 1 & Type 2 mezzanine adapter form factors, designed to mate into OCP servers
- OCP Spec 3.0 Small Form Factor (SFF), designed to mate into OCP servers
- Standard PCI Express Gen 3.0 and Gen 4.0 adapter cards
- Mellanox Socket Direct[®] adapter cards



OCP2.0 Adapter Card



OCP3.0 Adapter Card



Standard PCI Express Stand-up Adapter Card



Mellanox Socket Direct Adapter Card