




NVIDIA MELLANOX BLUEFIELD-2 DATA PROCESSING UNIT (DPU)



NVIDIA® Mellanox® BlueField®-2 is a highly-integrated Data Processing Unit (DPU) delivering advanced functionality, unmatched performance and agility for today's most demanding workloads. By combining the industry leading ConnectX®-6 Dx network adapter with an array of Arm® cores, BlueField-2 delivers a perfect blend of hardware and software accelerations.

Whether employed as a SmartNIC or as an embedded Storage Controller, BlueField-2 brings agile and high-performance solutions for security, machine learning, cloud, edge computing, and storage applications while reducing total cost of ownership (TCO).

BlueField-2 is the second generation in the BlueField family of DPUs. It integrates an array of eight powerful 64-bit Armv8 A72 cores interconnected by a coherent mesh with a DDR4 memory controller and a dual-port Ethernet or InfiniBand network adapter.

BlueField-2 supports two ports of 10/25/50/100 Gb/s or one port of 200 Gb/s Ethernet or InfiniBand, one out-of-band management port for the Arm subsystem, and includes an integrated 16-lane PCIe Gen 3.0/4.0 switch, offering endpoint and root complex capabilities.

At the heart of BlueField-2 is the industry leading ConnectX-6 Dx network adapter with unprecedented built-in hardware functionality, performance and resiliency. BlueField-2 hardware accelerations include advanced RDMA/RoCE capabilities as well as innovative cryptography, storage and networking accelerations. Relying on these built-in offloads, the BlueField-2 powerful array of Arm cores can be programmed for sophisticated custom accelerations and control path manipulations. BlueField-2 also benefits from the rich Arm ecosystem to deliver complete solutions to cloud operators and service providers.

Key Applications

- > Virtualized and bare metal cloud services
- > NVMe storage virtualization
- > Network Function Virtualization (NFV)
- > Security applications such as DDoS and Deep Packet Inspection (DPI)
- > Microservers built for edge computing

Key Features

- > Eight 64-bit Armv8 A72 cores
- > Dual ports of 10/25/50/100 Gb/s or a single port of 200 Gb/s
- > Supports Ethernet and InfiniBand protocols
- > DDR4 memory controller with ECC support
- > 16-lane PCIe Gen3.0/4.0 switch
- > Software programmability
- > Hardware offloads including:
 - > ASAP² - Accelerated Switch and Packet Processing[®]
 - > NVMe SNAP[™] storage emulation
 - > IPsec/TLS data-in-motion and AES-XTS data-at-rest crypto accelerations
 - > Regular expression (RegEx) and DPI accelerations
 - > SHA 256-bit hardware acceleration
- > Hardware Root of Trust
- > Host isolation
- > Zero-Touch RoCE (ZTR)
- > Supports up to 4 multi-hosts
- > 1GbE out-of-band management port

TURNING ZERO TRUST TO HERO TRUST

BlueField-2 delivers isolation, security and cryptography accelerations, enabling data center security at the endpoint with best-in-class performance, turning zero-trust to hero-trust.

Isolation – BlueField-2 as a SmartNIC functions as a “computer-in-front-of-a-computer”, unlocking unlimited opportunities for custom security applications on its Arm processors, fully isolated from the host’s CPU. In the event of a compromised host, BlueField-2 may detect and block malicious activities in real time and at wire speed to prevent the attack from spreading further.

Cryptography Accelerations – From IPsec and TLS data-in-motion inline encryption to AES-XTS block-level data-at-rest encryption and public key acceleration, BlueField-2 hardware-based accelerations offload the crypto operations and free up the CPU, reducing latency and enabling scalable crypto solutions. BlueField-2 “host-unaware” solutions may transmit and receive data, while BlueField-2 acts as a bump-in-the-wire for crypto.

Securing Workloads – BlueField-2 accelerates connection tracking with its ASAP² - Accelerated Switch and Packet Processing[®] technology to enable stateful filtering on a per connection basis. Moreover, BlueField-2 includes a Titan IC regular expression (Titan RXP™) acceleration engine supported by IDS/IPS tools to detect host introspection and application recognition in real time (NBAR).

Secure NIC – BlueField-2 hardware leverages Root-of-Trust secured boot and Arm TrustZone technology to ensure integrity of the firmware and hardware. Using hardware keys and a trusted execution environment, BlueField-2 can serve as the keys’ safebox providing anti-cloning and zero-trust access solutions.

EFFICIENT AND HIGH-PERFORMANCE CLOUD DEPLOYMENTS

Virtualized Cloud – By leveraging BlueField-2 virtualization offloads, data center administrators can benefit from better server utilization, allowing more virtual machines and more tenants on the same hardware, while reducing the TCO and power consumption.

BlueField-2 virtualization accelerations include:

- > Mellanox ASAP² for Open vSwitch (OVS) with flexible, highly-efficient virtual switching and routing capabilities. OVS accelerations can be further customized using BlueField-2 Arm processing power to offload the control path.
- > Network overlay technology (VXLAN, NVGRE, Geneve) offloads, including encapsulation and decapsulation, over tunneled protocols and Network Address Translation (NAT) routing capabilities.

Bare Metal Deployments – BlueField-2 in bare metal deployments presents several distinct advantages. For instance, it enables cloud providers to provision servers while retaining the ability to control them and deploy policy rules in an isolated and secured manner, totally transparent to the tenant.

More Smarts in the Edge – The BlueField-2 broad set of capabilities delivered in a compact design is perfectly positioned for delivering microserver applications at the edge of the network.

A COMPLETE SOLUTION FOR STORAGE DEPLOYMENTS

Complete End-to-End Storage Solution – BlueField-2 unique storage and NVMe capabilities may be leveraged on both the storage target and the storage client. In high-performance storage arrays, BlueField-2 functions as the system’s main CPU, handling storage target tasks and traffic termination. As a SmartNIC, it may operate as a co-processor to offload specific storage tasks from the host.


Best in Class Storage Performance – Leveraging unique RDMA/RoCE capabilities, BlueField-2 delivers industry leading NVMe-oF performance, coupled with advanced built-in hardware storage offloads including compression/decompression, data deduplication, and signature handover T10-DIF.

Virtualize Your Storage with NVMe SNAP™ – NVMe SNAP logically presents networked storage as a local NVMe drive on the PCIe bus to host software, thus bringing virtualized storage to bare-metal clouds, making composable storage simple, while promoting CAPEX and OPEX savings.

SOFTWARE ENVIRONMENT

The NVIDIA Mellanox software package enables developers to fully utilize the BlueField-2 rich feature set and provides the following capabilities.

- > Fully compatible with the first generation BlueField
- > Leverages the growing Arm ecosystem to create diverse solutions and tune the applications for maximum performance
- > Ubuntu/CentOS commercial distribution
- > Yocto Software Development Kit (SDK) with open API to enable customized development
- > Out-of-band 1GbE management interface



ORDERING INFORMATION

OPN	Description
M42M08T22A0-NDTTEV	BlueField®-2 E-Series DPU, 8 Cores, Dual 100GbE/EDR VPI Ports, PCIe 4.0 x16, No Crypto
M42M08T22A0-CDTTEV	BlueField®-2 E-Series DPU, 8 Cores, Dual 100GbE/EDR VPI Ports, PCIe 4.0 x16, Crypto
M42M08P22A0-NDTTEV	BlueField®-2 P-Series DPU, 8 Cores, Dual 100GbE/EDR VPI Ports, PCIe 4.0 x16, No Crypto
M42M08P22A0-CDTTEV	BlueField®-2 P-Series DPU, 8 Cores, Dual 100GbE/EDR VPI Ports, PCIe 4.0 x16, Crypto
M42M08P12A0-NDTTHV	BlueField®-2 P-Series DPU, 8 Cores, Single 200GbE/HDR VPI Port, PCIe 4.0 x16, No Crypto
M42M08P12A0-CDTTHV	BlueField®-2 P-Series DPU, 8 Cores, Single 200GbE/HDR VPI Port, PCIe 4.0 x16, Crypto

E-Series - Efficient Performance | P-Series - High Performance.

FEATURES*

Network and Host Interfaces

Network Interfaces

- > Ethernet:
 - > Dual ports of 10/25/50/100 GbE, or
 - > Single port of 200 GbE
- > InfiniBand:
 - > Dual ports of up to EDR/HDR100 (100 Gb/s), or
 - > Single port of HDR (200 Gb/s)

PCI Express Interface

- > 16 lanes of PCIe Gen 3.0/4.0
- > Endpoint or root complex
- > PCIe switch with up to 8 downstream ports
- > Support for MSI/MSI-X

Mellanox Multi-Host

- > Mellanox Multi-Host® for up to 4 hosts

Arm/DDR Subsystem

Arm Cores

- > Up to 8 Armv8 A72 cores (64-bit) pipeline
- > Arm NEON™ 128b SIMD execution unit
- > Arm VFPv4 single and double precision floating point acceleration (IEEE 754)
- > Cache coherent mesh interconnect
- > Each two Arm cores share 1 MB L2 cache
- > 6MB L3 cache with plurality of eviction policies

DDR4 DIMM Support

- > Single DDR4 DRAM controller
- > Persistent memory support
- > Up to 256GB per channel
- > ECC error protection support

Hardware Accelerations

Security

- > IPsec/TLS data-in-motion encryption
 - > AES-GCM 128/256 bit key
- > AES-XTS 256/512 bit data-at-rest encryption
- > SHA 256 bit hardware acceleration
- > Regular expression (RegEx) acceleration

- > Arm A64, A32 & T32 cryptography instructions for:
 - > AES, SHA-1, SHA-224, and SHA-256
 - > Finite field arithmetic
- > Hardware Public Key Accelerator
 - > RSA, Diffie-Hellman, DSA, ECC, EC-DSA, EC-DH
- > True Random Number Generator (TRNG)
- > Hardware Root of Trust
 - > Cerberus compliant

Storage

- > NVMe SNAP™
- > NVMe over Fabric offloads for target
- > T10-DIF signature handover
- > Dedup accelerations
- > DMA accelerations
- > Compression/decompression acceleration

RDMA/RoCE

- > Zero-touch RoCE configuration
- > Selective repeat
- > Programmable congestion control
- > Collective / vector collective operations
- > GPUDirect®
- > Enhanced atomic operations
- > 16 million I/O channels

TCP/IP Transport

- > Stateless offloads for:
 - > TCP/UDP/IP
 - > LSO/LRO/Checksum/RSS/TSS/HDS
 - > VLAN insertion/stripping

Hardware-based I/O Virtualization

- > SR-IOV
 - > Up to 1K Virtual Functions
 - > Up to 56 Physical Functions per host
- > VirtIO
- > Multi-function per port
- > VMware NetQueue support
- > Virtualization hierarchies
- > Virtualizing physical functions on a physical port
- > 1K ingress and egress QoS levels
- > Guaranteed QoS for VMs

Overlay Networks

- > Hardware offload of encapsulation and decapsulation of NVGRE, VXLAN and Geneve
- > Header rewrite (NAT)

Management and Boot

Advanced Boot Options

- > Secure boot (RSA authenticated)
- > Remote boot over Ethernet/InfiniBand
- > Remote boot over iSCSI
- > PXE and UEFI

Management and Control Interfaces

- > NC-SI, MCTP over SMBus and MCTP over PCIe
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP026
- > 1GbE out-of-band management port
- > SDN management interface for managing the eSwitch
- > I²C interface for device control and configuration
- > General Purpose I/O pins
- > SPI interface to flash
- > eMMC memory controller
- > MDC/MDIO master
- > UART
- > USB
- > JTAG IEEE 1149.1 and IEEE 1149.6

Software

Software Development Toolchain

- > Native and cross-compile GNU toolchain
- > Compatible with Arm DS-5 and other commercial development and profiling tools

Software Support

- > Ubuntu/CentOS Commercial Linux Distribution
- > Yocto-based Linux Distribution
- > OpenFabrics Enterprise Distribution (OFED)
- > Arm-optimized versions of all Mellanox drivers and software stack
- > IDS/IPS tools such as Snort and Suricata
- > Optimized Arm DPDK and ConnectX PMD

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

Learn more at www.mellanox.com/products/bluefield2-overview