



NVIDIA® Mellanox® BlueField® SmartNIC for InfiniBand & Ethernet

High Performance Smart Network Adapter Cards

Combining Arm® processing power with advanced network and storage offloads to accelerate a multitude of security, networking and storage applications for cloud computing and high performance computing (HPC) with world-leading performance, flexibility and efficiency.

BlueField SmartNIC features the BlueField Data Processing Unit (DPU) — an innovative and high-performance programmable networking engine. Providing unmatched scalability and efficiency, the dual-port BlueField SmartNIC is the ideal adapter to accelerate the most demanding workloads.

Powerful & Smart Network Adapter

BlueField SmartNIC adapters accelerate a wide range of applications through flexible data and control plane offloading. Enabling a more efficient use of compute resources, BlueField adapters empower the CPU to focus on running applications rather than on networking or security processing. Additionally, as software-defined adapters, BlueField SmartNICs ensure the ultimate flexibility by adapting to future protocols and features through simple software updates.

BlueField System-on-Chip

BlueField is a Mellanox family of advanced solutions that integrates a coherent mesh of 64-bit Armv8 A72 cores, a ConnectX network adapter front-end and a PCI Express switch into a single chip.

The powerful DPU architecture includes a 64-bit Armv8 multicore processor array and enables customers to develop sophisticated applications and highly differentiated feature sets. BlueField leverages the rich Arm software ecosystem and introduces the ability to offload the x86 software stack.

At the heart of BlueField is the ConnectX-5 network controller with DMA over Converged Ethernet (RoCE) and InfiniBand offload technology, which delivers cutting-edge performance for networking and storage applications such as NVMe over Fabrics. Advanced features include an embedded virtual switch with programmable access lists (ACLs), transport offloads and stateless protocols.

ConnectX-5 offers enhancements to HPC infrastructures, by providing MPI and SHMEM/PGAS and Rendezvous Tag Matching offloads, hardware support for out-of-order RDMA Write and Read operations, as well as additional network atomic operation (NAO) support.



HIGHLIGHTS

- Intelligent programmable network adapter
- Best-in-class hardware offloads with Arm processing power
- Accelerates a wide range of security, networking, storage and other workloads
- Dedicated hardware offload for NVMe-over-Fabrics (NVMe-oF)
- Dual-port 100Gb/s PCIe Gen4.0 x16
- Supports InfiniBand and Ethernet protocols
- 4/8/16GB on-board DDR4 memory, depending upon model
- FHHL form factor
- Standard embedded Linux software stack
- Hardware-based security isolation

Security Applications

BlueField SmartNIC addresses the concerns of modern data centers by combining hardware encryption accelerators with embedded software and fully integrated advanced network capabilities, making it an ideal platform for developing proprietary security applications.

It enables a distributed security architecture by isolating and protecting each individual workload and providing flexible control and visibility at the server and workload level, controlling risk at the server access layer. BlueField builds security into the DNA of the data center and enables prevention, detection and response to potential threats in real time.

BlueField SmartNIC is capable of delivering powerful functionality, including encryption of data-in-motion, bare metal provisioning, stateful L4 firewall and more.

Networking Applications

BlueField SmartNIC offers a wide range of dedicated offloads to maximize virtualization scalability and efficiency. Data center administrators can benefit from better server utilization, allowing more virtual machines and more tenants on the same hardware, while reducing the TCO, power, and cable complexity.

Among its accelerations:

- Mellanox ASAP² - Accelerated Switching and Packet Processing[®] for Open vSwitch (OVS) delivers flexible, highly efficient virtual switching and routing capabilities. OVS accelerations can be further enhanced using BlueField processing and memory. For example, the scale of OVS actions can be increased by utilizing BlueField internal memory, and more OVS actions and vSwitch/vRouter implementations can be supported.
- Network overlay technology (VXLAN, NVGRE, Geneve) offload, including encapsulation and decapsulation, allows the traditional offloads to operate on the tunneled protocols and also offload Network Address Translation (NAT) routing capabilities.

Storage Applications

BlueField SmartNIC may operate as a co-processor offloading specific storage tasks from the host, isolating part of the storage media from the host, or enabling abstraction of software-defined storage logic, using both a BlueField-dedicated hardware offload for NVMe-over-Fabrics and Arm cores.

On the storage initiator side, BlueField SmartNICs can prove an efficient solution for hyperconverged systems to enable the host CPU to focus on compute while all the storage and networking tasks are handled through the NIC engine and Arm cores.

Decoupling of the storage tasks from the compute tasks also simplifies the software model, enabling the deployment of multiple OS virtual machines while the storage application is handled solely by the Arm Linux subsystem.

Mellanox NVMe SNAP[™]

NVMe SNAP (Software-defined Network Accelerated Processing) enables hardware virtualization of NVMe storage. The Mellanox NVMe SNAP framework enables customers to easily integrate networked storage solutions into their cloud or enterprise server deployments. NVMe SNAP brings virtualized storage to bare-metal clouds and makes composable storage simple, promoting CAPEX and OPEX savings. It enables the efficient disaggregation of compute and storage to allow fully-optimized resource utilization.

NVMe SNAP logically presents networked storage as a local NVMe drive on the PCIe bus to host software. This allows the host OS/Hypervisor to use a standard NVMe-driver instead of different network storage protocols.

NVMe SNAP empowers customers with the freedom to implement their own storage technology and solutions on top of the NVMe SNAP framework that runs on Mellanox BlueField embedded Arm cores along with embedded hardware storage acceleration engines. This powerful combination is agile yet completely transparent to host software, allowing NVMe SNAP to be integrated into almost any storage solution.

NVMe SNAP is a licensed software available on select BlueField part numbers, see Table 1 for more information.

Software Support

BlueField SmartNIC is shipped with Mellanox BlueOS[®] but other operating systems, such as RH xx.xx and others, are also supported. BlueOS is a Linux reference distribution, which includes the Mellanox OFED stack, and is capable of running all customer-based Linux applications seamlessly. BlueOS itself is based on the Yocto Project Poky distribution. A PXE driver is also pre-installed.

The BlueField adapter execution environment is fully isolated from the x86 server and network environment, and can run the Open vSwitch Database (OVSDb) or other virtual switches to create a secure solution for bare metal provisioning.

The software package also includes support for DPDK, and applications for encryption and a stateful L4-based firewall.



Features

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

NETWORK AND HOST INTERFACES

Network Interfaces

- Dual-port QSFP28
- Integrated PHYs seamlessly connect to all standard copper and fiber media

PCI Express Interface

- 8 or 16 lanes of PCIe Gen 3.0/4.0
- Fall-back to 4, 2, or 1 lane

ARM/DDR SUBSYSTEM

Arm Cores

- Armv8 A72 cores (64-bit)
- Superscalar, variable-length, out-of-order pipeline
- Each core supports NEON™ 128b SIMD execution unit
- Arm VFPv4 single and double precision floating point acceleration (IEEE 754)
- Per core 48KB I-cache and 32KB D-cache
- Cache coherent mesh interconnect of CPUs, I/O and memory – each tile contains 2 cores and 1 MB L2 cache
- 6 MB L3 cache, sophisticated eviction policies

On-board Memory

- Single channel 4/8/16GB DDR4 DRAM with ECC

HARDWARE ACCELERATIONS

Encryption Acceleration

- Armv8 cryptography extensions: A64, A32, and T32 instructions for:
 - AES, SHA-1, SHA-224, and SHA-256
 - Finite field arithmetic used in algorithms such as Galois/Counter Mode and Elliptic Curve
- Hardware Public Key Accelerator
 - RSA, Diffie-Hellman, DSA, ECC, EC-DSA, EC-DH
- True Random Number Generator with entropy source

Enhanced Features

- PeerDirect RDMA (aka GPUDirect) communication acceleration
- Enhanced Atomic operations
- Registration-free RDMA memory access
- NVMe-over-Fabric hardware acceleration

Transport Offloads

- RDMA over Converged Ethernet (RoCE) and InfiniBand
- TCP/UDP/IP stateless offload
- LSO, LRO, checksum offload
- RSS (also on encapsulated packets), TSS, HDS, VLAN insertion/stripping, Receive Flow Steering
- Intelligent interrupt coalescence
- TCP/UDP, MPLS, VxLAN, NVGRE, GENEVE
- SRP, iSER, NFS RDMA, SMB Direct, NVMe-oF

Hardware-based I/O Virtualization

- SR-IOV
- Multi-function per port
- Multiple queues per virtual machine

Management and Control

- SDN management interface for managing the eSwitch
- eMMC memory controllers

SOFTWARE DEVELOPMENT TOOLCHAIN

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

SOFTWARE SUPPORT

Arm Environment

- BlueOS: Commercial grade Yocto-based Arm Linux distribution
- Commercial Linux distributions supported
- Delivered with OpenFabrics Enterprise Distribution (OFED)
- Arm-optimized versions of all Mellanox drivers and software stack
- Accelerated NVMe over Fabrics target stack
- Optimized Arm DPDK and ConnectX PMD
- NVMe SNAP framework (license required)

Connected Host (Network Adapter Environment)

- Linux
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF-2)

Table 1 - Part Numbers and Feature Set Breakdown

OPN	Max. Speed*	No. of Ports	PCIe Support	Cores	NVMe SNAP™	Crypto**	DDR Memory	Form Factor
MBF1L516A-ESNAT	100Gb/s	2 x QSFP28	Gen4.0 x16	G-Series 16 Cores	Optional - Contact Mellanox	Crypto disabled	16GB on-board	FHHL Tall Bracket
MBF1L516A-ESCAT	100Gb/s	2 x QSFP28	Gen4.0 x16			Crypto enabled	16GB on-board	

Notes: *Refer to BlueField Software and Firmware release notes for availability of FDR 56Gb/s capabilities.

**Secure boot is not supported on these SmartNIC models, please contact Mellanox for additional information.

Support

For information about Mellanox support packages, please contact your Mellanox Technologies sales representative or visit our [Support Index page](#).