




NVIDIA BLUEFIELD-2 DPU

DATA CENTER INFRASTRUCTURE ON A CHIP



The NVIDIA® BlueField-2® data processing unit (DPU) is a data center infrastructure on a chip optimized for traditional enterprise, high-performance computing (HPC), and modern cloud workloads, delivering a broad set of accelerated software-defined networking, storage, security, and management services. With its 200 Gb/s InfiniBand or Ethernet connectivity, the BlueField-2 DPU enables organizations to transform their IT infrastructures into state-of-the-art data centers that are accelerated, fully programmable, and armed with “zero trust” security to prevent data breaches and cyber attacks.

By combining the industry-leading NVIDIA Mellanox® ConnectX-6® Dx network adapter with an array of Arm® cores, BlueField-2 offers purpose-built, hardware-acceleration engines with full software programmability. Sitting at the edge of every server, BlueField-2 is optimized to handle critical infrastructure tasks quickly, increasing data center efficiency.

BlueField-2 empowers agile and high-performance solutions for cloud networking, storage, cyber security, data analytics, HPC, and artificial intelligence (AI), from edge to core data centers and clouds, all while reducing the total cost of ownership.

The NVIDIA DOCA software development kit (SDK) enables developers to rapidly create applications and services for the BlueField-2 DPU. The DOCA SDK makes it easy and straightforward to leverage DPU hardware accelerators and CPU programmability for better application performance and security.

KEY SOFTWARE-DEFINED, HARDWARE-ACCELERATED APPLICATIONS

Networking

- > vSwitch/vRouter, NAT, load balancer, NFV

Storage

- > NVMe over Fabric (NVMe-oF) Storage Direct, encryption, elastic storage, data integrity, compression, data deduplication

Security

- > Next-Generation firewall, IDS/IPS, root of rust, micro-segmentation, DDOS prevention

PORTFOLIO

- > Dual ports of up to 100 Gb/s, or a single port of 200 Gb/s
- > 8 GB / 16 GB of on-board DDR4 memory
- > Form factors: HHHL, FHHL, and OCP 3.0 SFF
- > 1 GbE out-of-band management port

KEY FEATURES

Security

- > Hardware root of trust
- > Regular expression (RegEx) matching processor acceleration
- > Hardened isolation layer
- > IPsec/TLS and XTS encryption
- > Connection tracking for stateful firewall acceleration

Storage


- > GPUdirect® Storage
- > Elastic block storage enabled by NVMe SNAP™ storage virtualization
- > NVMe-oF accelerator
- > VirtIO-blk accelerator

Networking

- > RoCE, Zero Touch RoCE
- > GPUdirect
- > VirtIO-net accelerator
- > Accelerated Switching and Packet Processing® (ASAP²)
- > OVS and OVN acceleration

Management

- > Authenticated product life-cycle management
- > Telemetry agents



BlueField-2 Silicon Architecture

FEATURES

NETWORK AND HOST INTERFACES

Network Interfaces

- > Ethernet - Dual ports of 10/25/50/100 Gb/s, or a single port of 200 Gb/s
- > InfiniBand - Dual ports of EDR/HDR100 (100 Gb/s), or a single port of HDR (200 Gb/s)

PCI Express Interface

- > 8 or 16 lanes of PCIe Gen 4.0
- > PCIe switch with up to 8 downstream ports

ARM/DDR SUBSYSTEM

Arm Cores

- > Up to 8 Armv8 A72 cores (64-bit) pipeline
- > Arm NEON™ 128b SIMD execution unit
- > 1 MB L2 cache per 2 cores
- > 6 MB L3 cache with plurality of eviction policies

DDR4 DIMM Support

- > Single DDR4 DRAM controller
- > 8 GB / 16 GB on-board DDR4
- > ECC error protection support

HARDWARE ACCELERATIONS

Security

- > Hardware root of trust
 - > Cerberus compliant
- > Regular expression (RegEx) acceleration
- > 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
- > Hardware public key accelerator
 - > RSA, Diffie-Hellman, DSA, ECC, EC-DSA, EC-DH
- > True random number generator (TRNG)

Storage

- > NVMe SNAP
- > VirtIO-blk
- > NVMe-oF acceleration
- > NVMe/TCP
- > Compression and decompression
- > Data hashing and deduplication

Network

- > RoCE, Zero Touch RoCE
- > Stateless offloads
 - > TCP/UDP/IP
 - > LSO/LRO/Checksum/RSS/TSS/HDS
 - > VLAN insertion/stripping
- > SR-IOV
 - > Up to 1K virtual functions
 - > Up to 56 physical functions per host
- > VirtIO-net
- > Multi-function per port
- > VMware NetQueue support
- > Virtualization hierarchies
- > 1K ingress and egress QoS levels
- > Accelerations for tunneling technologies: NVGRE, VXLAN, Geneve
- > Header rewrite (NAT)

AI Accelerations

- > GPUDirect
- > GDS (GPUDirect Storage)

BOOT OPTIONS

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

HOST MANAGEMENT

- > 1 GbE out-of-band management port
- > NC-SI, MCTP over SMBus, and MCTP over PCIe
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP026
- > I²C interface for device control and configuration
- > SPI interface to flash
- > eMMC memory controller > UART
- > USB


SOFTWARE

DOCA Software Development Kit


- > Network acceleration SDK: ASAP² SDN, emulated VirtIO, 5T PTP technology
- > Security acceleration SDK: Inline cryptography, deep packet inspection
- > Storage acceleration SDK: Storage emulation and virtualization, crypto, and compression
- > RDMA acceleration SDK: UCX, RDMA verbs, GPUDirect
- > I/O programmability SDK: P4 and P4RT, OpenSNAPI, cuNE

OS Support


- > VMware Cloud Foundation
- > CentOS
- > Ubuntu Commercial Linux Distribution
- > RedHat
- > SuSE
- > Windows
- > Debian 9.11
- > Yocto-based Linux Distribution
 - > OpenFabrics Enterprise Distribution (OFED)
 - > Arm-optimized versions of all NVIDIA drivers and software stack




BlueField-2 DPU - 2x 25 Gb/s HHHL form factor



BlueField-2 DPU - 2x 100 Gb/s FHHL form factor



BlueField-2 DPU - 2x 25 Gb/s OCP3.0 SFF form factor



BlueField-2 Silicon

ORDERING INFORMATION

PART NUMBERS AND FEATURE SET BREAKDOWN^{1,2,3}

OPN	Max. Speed	No. of Ports	PCIe Support	Cores Speed	Crypto	DDR Memory	1GbE OOB	Form Factor
MBF2M322A-AEEOT	25 GbE 2x	SFP56	Gen 4.0 x8	2.0 GHz	Crypto enabled	8 GB on-board	Yes	HHL Tall Bracket
MBF2H322A-AEEOT	25 GbE 2x	SFP56	Gen 4.0 x8	2.5 GHz	Crypto enabled	8 GB on-board	Yes	HHL Tall Bracket
MBF2M912A-AEEAB	25 GbE 2x	SFP56	Gen 3.0 x16	2.0 GHz	Crypto enabled	8 GB on-board	No	OCP 3.0
MBF2M516A-CEEOT	100 GbE 2x	QSFP56	Gen 4.0 x16	2.0 GHz	Crypto enabled	16 GB on-board	Yes	FHL Tall Bracket
MBF2H516A-CEEOT	100 GbE 2x	QSFP56	Gen 4.0 x16	2.5 GHz	Crypto enabled	16 GB on-board	Yes	FHL Tall Bracket
MBF2M516A-EEEOT	EDR/HDR100 2x	QSFP56	Gen 4.0 x16	2.0 GHz	Crypto enabled	16 GB on-board	Yes	FHL Tall Bracket
MBF2H516A-EEEOT	EDR/HDR100 2x	QSFP56	Gen 4.0 x16	2.5 GHz	Crypto enabled	16 GB on-board	Yes	FHL Tall Bracket

1 Hardware root of trust (RoT) or secure boot is not supported on these models. Please contact NVIDIA for additional information.

2 All SKUs are also available with crypto-disabled mode. Please contact NVIDIA for additional information.

3 All SKUs are available with 8 GB or 16 GB DDR4. Please contact NVIDIA for additional information.

For silicon SKUs, please contact the NVIDIA sales team.

Support:

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

LICENSES SOLD SEPARATELY

OPN	Description
BF1-NVMESNAP-BNS-1	One perpetual license to use NVMe SNAP on one adapter of BlueField-2. Includes Mellanox Technical Support and Warranty – Silver, 1 Year.
Contact Mellanox	One perpetual license to use RegEx acceleration on one adapter of BlueField-2. Includes Mellanox Technical Support and Warranty – Silver, 1 Year.

Learn more at www.nvidia.com/DPU

© 2020 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, Accelerated Switch and Packet Processing (ASAP2), BlueField, ConnectX, GPUDirect, Mellanox, and NVMe SNAP are trademarks and/or registered trademarks of NVIDIA Corporation 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. ARM, AMBA, and ARM Powered are registered trademarks of ARM Limited. Cortex, MPCore and Mali are trademarks of ARM Limited. "ARM" is used to represent ARM Holdings plc; its operating company ARM Limited; and the regional subsidiaries ARM Inc.; ARM KK; ARM Korea Limited.; ARM Taiwan Limited; ARM France SAS; ARM Consulting (Shanghai) Co. Ltd.; ARM Germany GmbH; ARM Embedded Technologies Pvt. Ltd.; ARM Norway, AS and ARM Sweden AB. OCT20

