


# NVIDIA BLUEFIELD-2 DPU

## DATA CENTER INFRASTRUCTURE ON A CHIP



The NVIDIA® BlueField®-2 data processing unit (DPU) is the world's first data center infrastructure-on-a-chip optimized for traditional enterprises' modern cloud workloads and high performance computing. It delivers a broad set of accelerated software-defined networking, storage, security, and management services with the ability to offload, accelerate and isolate data center infrastructure. With its 200Gb/s Ethernet or InfiniBand 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 ConnectX®-6 Dx network adapter with an array of Arm® cores and infrastructure-specific offloads, BlueField-2 offers purpose-built, hardware-acceleration engines with full software programmability. Sitting at the edge of every server, BlueField-2 empowers agile, secured and high-performance cloud and artificial intelligence (AI) workloads, all while reducing the total cost of ownership and increasing data center efficiency.

The NVIDIA DOCA™ software framework enables developers to rapidly create applications and services for the BlueField-2 DPU. NVIDIA DOCA makes it easy to leverage DPU hardware accelerators, providing breakthrough data center performance, efficiency and security.

## KEY SOFTWARE-DEFINED, HARDWARE-ACCELERATED APPLICATIONS



### Networking

vSwitch/vRouter, NAT, load balancer, NFV



### Storage

NVMe™ over fabrics (NVMe-oF™), elastic storage virtualization, hyper converged infrastructure (HCI), encryption, data integrity, compression, data deduplication



### Security

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

## Key Features

### Security

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

### Storage

- > NVIDIA GPUDirect® Storage
- > Elastic block storage enabled by BlueField SNAP storage virtualization
- > Compression and decompression acceleration
- > NVMe-oF acceleration
- > VirtIO-blk acceleration

### Networking

- > RoCE, Zero Touch RoCE
- > GPUDirect
- > SDN acceleration powered by NVIDIA ASAP² - Accelerated Switching and Packet Processing®
- > Overlay network offloads including VXLAN

### Management

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

## Portfolio

- > Dual ports of up to 100Gb/s, or a single port of 200Gb/s Ethernet or InfiniBand
- > 8GB / 16GB / 32GB of on-board DDR4 memory
- > Card form factors: HHHH, FHHL, and OCP 3.0 SFF
- > M.2 / U.2 connectors for direct attached storage
- > 1GbE out-of-band management port

# FEATURES

## Network and Host Interfaces

### Network Interfaces

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

### PCI Express Interface

- > 8 or 16 lanes of PCIe Gen 4.0
- > PCIe switch bi-furcation with 8 downstream ports

## ARM/DDR Subsystem

### Arm Cores

- > Up to 8 Armv8 A72 cores (64-bit) pipeline
- > 1MB L2 cache per 2 cores
- > 6MB L3 cache with plurality of eviction policies

### DDR4 DIMM Support

- > Single DDR4 DRAM controller
- > 8GB / 16GB / 32GB of on-board DDR4
- > ECC error protection support

## Hardware Accelerations

### Security

- > Secure boot with hardware root-of-trust
  - > Secure firmware update
  - > 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

- > BlueField SNAP - NVMe™ and VirtIO-blk
- > NVMe-oF™ acceleration
- > Compression and decompression acceleration
- > Data hashing and deduplication
- > M.2 / U.2 connectors for direct attached storage

### Networking


- > RoCE, Zero Touch RoCE
- > Stateless offloads for:
  - > TCP/UDP/IP
  - > LSO/LRO/checksum/RSS/TSS/HDS
  - > VLAN insertion/stripping
- > SR-IOV
- > VirtIO-net
- > Multi-function per port
- > VMware NetQueue support
- > Virtualization hierarchies
- > 1K ingress and egress QoS levels

### Boot Options


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

### Management


- > 1GbE 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



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



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



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

## ORDERING INFORMATION

For information about NVIDIA ordering information, please contact your NVIDIA sales representative or visit our [Nvidia BlueField-2 User Guide index page](#):

**NVIDIA BlueField-2 Ethernet boards**

**NVIDIA BlueField-2 InfiniBand/VPI boards**

**NVIDIA BlueField-2 for OCP3.0**

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

To learn more about the NVIDIA BlueField-2 visit [www.nvidia.com/dpu](http://www.nvidia.com/dpu)

© 2021 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, Accelerated Switch and Packet Processing (ASAP2), BlueField, ConnectX, GPUDirect, Mellanox, and BlueField 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. APR21

