Datasheet

NVIDIA BlueField-3 Networking Platform

The 400Gb/s infrastructure compute platform for powering the world's data centers.

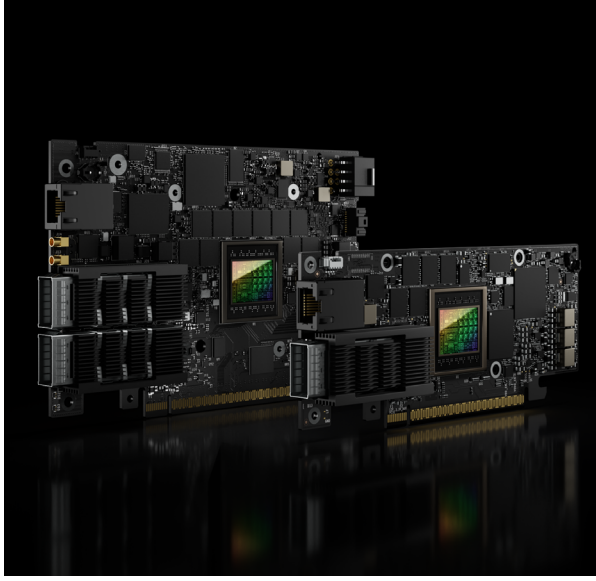

The NVIDIA® BlueField®-3 networking platform is designed to accelerate data center infrastructure workloads and usher in the era of accelerated computing and AI. Supporting both Ethernet and InfiniBand connectivity, BlueField-3 offers speeds up to 400 gigabits per second (Gb/s). It combines powerful computing with software-defined hardware accelerators for networking, storage, and cybersecurity—all fully programmable through the NVIDIA DOCA™ software framework. Drawing on the platform's robust capabilities, BlueField data processing units (DPUs) and BlueField SuperNICs revolutionize traditional computing environments, transforming them into secure, high-performance, efficient, and sustainable data centers suitable for any workload at any scale.

The **BlueField-3 DPU** is a cloud infrastructure processor that empowers organizations to build software-defined, hardware-accelerated data centers from the cloud to the edge. BlueField-3 DPUs offload, accelerate, and isolate software-defined networking, storage, security, and management functions, significantly enhancing data center performance, efficiency, and security. By decoupling data center infrastructure from business applications, BlueField-3 creates a secure, zero-trust data center infrastructure, streamlines operations, and reduces the total cost of ownership.

The **BlueField-3 SuperNIC** is a novel class of network accelerator that's purpose-built for supercharging hyperscale AI workloads. Designed for network-intensive, massively parallel computing, the BlueField-3 SuperNIC provides best-in-class remote direct-memory access over converged Ethernet (RoCE) network connectivity between GPU servers at up to 400Gb/s, optimizing peak AI workload efficiency. For modern AI clouds, the BlueField-3 SuperNIC enables secure multi-tenancy while ensuring deterministic performance and performance isolation between tenant jobs.

Portfolio

- 1 or 2 ports with up to 400Gb/s connectivity
- 32GB on-board DDR5 memory
- Form factors: HHHL, FHHL
- 1GbE out-of-band management port
- Integrated BMC

Key Software-Defined, Hardware-Accelerated Applications

Cloud Networking

Cloud overlay, SDN acceleration, NAT, load balancer, NFV, video streaming

Storage

NVMe™ over Fabrics (NVMe-oF™), NVMe/TCP™, elastic storage, hyper converged infrastructure (HCI)

Security

Distributed next-generation firewall, root of trust, micro-segmentation, DDOS prevention

HPC / AI

AI cloud, secure multi-tenancy, cloud-native supercomputing, communication acceleration

Telco and Edge

Cloud RAN, virtualized edge gateways, VNF acceleration, edge microservers

Features

Network and Host Interfaces

Network Interfaces

- 1 or 2 ports with up to 400Gb/s Ethernet or NDR InfiniBand connectivity

PCI Express Interface

- 32 lanes of PCIe Gen 5.0
- Flexible PCIe switch supporting self-hosting and server-hosting

Compute and Memory

Arm CPU Cores

- Up to 16 Armv8.2+ A78 Hercules cores
- 8MB L2 cache
- 16MB LLC system cache

Programmable Datapath Accelerator

- 16 cores, 256 threads
- Programmability through DOCA
- Heavy multi-threading applications acceleration

DDR and SSD Support

- Dual DDR5 5600MT/s DRAM controllers
- 32GB on-board DDR5
- ECC error protection support
- 128GB on-board SSD

Hardware Accelerations

Security

- Platform security
  - Secure boot with hardware root-of-trust
  - Secure firmware update
  - On-board flash encryption
  - Device attestation
- Functional isolation layer
- IPsec/TLS/MACSec 128/256bit data-in-motion encryption
- PSP security protocol (PSP)
- AES-GCM 128/256bit key
- AES-XTS 256/512bit data-at-rest encryption
- Connection tracking for statefull firewall
- Public key accelerator (PKA)
- True random number generator (TRNG)

Storage

- BlueField SNAP - Elastic block storage - NVMe™ and VirtIO-blk
- NVMe-oF™ and NVMe/TCP™ acceleration
- Decompression engine
- Erasure coding for RAID implementation

Networking

- RoCE, Zero Touch RoCE
- ASAP² - Accelerated Switch and Packet Processing® for SDN and VNF acceleration
- Single Root I/O Virtualization (SR-IOV)
- VirtIO acceleration
- Overlay network acceleration
  - VXLAN, Geneve, NVGRE
- Programmable flexible parser: user-defined classification
- Connection tracking (L4 firewall)
- Flow mirroring, sampling and statistics
- Programmable congestion control (PCC)
- Stateless TCP offloads

HPC/AI Accelerations

- HPC / AI All-to-All engine
- NVIDIA GPUDirect
- NVIDIA GPUDirect Storage (GDS)
- HPC MPI Tag Matching

Advanced Timing and Synchronization

- IEEE 1588v2 (any profile)
- PTP hardware clock (PHC)
- Line rate hardware timestamp
- Time triggered scheduling
- Time-based SDN acceleration

Boot Options

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

Management

- Integrated BMC
- 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 for storing the system's BIOS
- UART debug interface
- USB connector to load operating system images

Ready to Get Started?

For ordering information, please contact your NVIDIA sales representative or visit the **NVIDIA BlueField-3 User Guide**

To learn more about the NVIDIA BlueField Networking Platform, visit **nvidia.com/dpu**

NVIDIA