



NVIDIA BlueField-3 Networking Platform User Guide

Table of Contents

1	Introduction.....	10
1.1	System Requirements	10
1.2	Package Contents	11
1.2.1	Card Package	11
1.2.2	Accessories Kit	11
1.2.3	Optional PCIe Auxiliary Card Package	11
1.3	Features and Benefits	12
2	BlueField-3 Administrator Quick Start Guide	16
2.1	Prerequisites for Initial BlueField-3 Deployment	16
2.2	First-time Installation Procedure	16
3	Supported Interfaces	17
3.1	BlueField-3 SuperNICs Layout and Interface Information.....	17
3.2	BlueField-3 DPUs Layout and Interface Information.....	19
3.3	Interfaces Detailed Description.....	20
3.3.1	System-on-Chip (SoC)	20
3.3.2	Networking Interface.....	21
3.3.3	Networking Ports LEDs Interface	21
3.3.4	PCI Express Interface.....	22
3.3.5	DDR5 SDRAM On-Board Memory.....	22
3.3.6	NC-SI Management Interface	23
3.3.7	UART Interface Connectivity	24
3.3.8	USB 4-pin RA Connector.....	24
3.3.9	1GbE OOB Management Interface	24
3.3.10	PPS IN/OUT Interface	25
3.3.11	External PCIe Power Supply Connector.....	25
3.3.12	Cabline CA-II Plus Connectors.....	27
3.3.13	Integrated BMC Interface	27
3.3.14	NVMe SSD Interface.....	27
3.3.15	RTC Battery	28
3.3.16	eMMC Interface.....	28
4	Pinouts Description	29
4.1	PCI Express Interface	29

4.2	External Power Supply Connector	31
4.3	NC-SI Management Interface	32
4.4	Cabline CA-II Plus Connectors Pinouts	34
4.4.1	Component Side.....	34
4.4.2	Print Side	36
5	Hardware Installation and PCIe Bifurcation	39
5.1	Safety Warnings	39
5.2	Installation Procedure Overview.....	40
5.3	System Requirements	40
5.3.1	Hardware Requirements	40
5.3.2	Airflow Requirements	41
5.3.3	Software Requirements	41
5.4	Safety Precautions	41
5.5	Unpacking	41
5.6	Pre-Installation Checklist	41
5.7	Installation Instructions	42
5.8	Cables and Modules.....	42
5.8.1	Networking Cable Installation	42
5.8.2	8-pin ATX Power Supply Cable	43
5.8.3	Cable Installation Instructions	44
5.9	DPU Power-Up and Power-Down Sequences	45
5.9.1	Power-Up Sequence	45
5.9.2	Power-Down Sequence	46
5.10	PCIe x16 Installation Instructions.....	47
5.10.1	Installation Instructions.....	47
5.10.2	Uninstalling the BlueField-3 Card	48
5.11	[Optional] PCIe Extension Connection (2x PCIe x16) Installation Instructions	49
5.11.1	Installing the DPU.....	50
5.11.2	Uninstalling the Cards.....	54
5.12	PCIe Bifurcation Configuration Options	55
5.12.1	Host as Root Port on x4 PCIe Lane Peripherals	57
5.12.2	DPU ARMs as Root Port on Peripherals	58
6	Setting High-Speed-Port Link Type	60
6.1	mlxconfig.....	60

6.2	UEFI	60
7	Troubleshooting	61
8	Specifications.....	62
8.1	B3140H SuperNICs Specifications	62
8.2	B3140L SuperNICs Specifications	63
8.3	B3220L SuperNICs Specifications	64
8.4	B3210L SuperNICs Specifications	66
8.5	B3240 DPUs Specifications.....	67
8.6	B3240 DPU for Cold-Aisle Servers Specifications.....	68
8.7	B3210 DPUs Specifications.....	69
8.8	B3210E DPUs Specifications	70
8.9	B3220 DPUs Specifications.....	71
8.10	DPUs Mechanical Drawing and Dimensions	73
8.11	Bracket Mechanical Drawings	74
9	Monitoring.....	76
9.1	Thermal Sensors.....	76
9.2	Heatsink	76
10	Finding the GUID/MAC on the Board	77
10.1	DPUs Board Label Example	77
10.2	SuperNICs Board Label Example	79
11	PCIe Auxiliary Card Kit	81
11.1	Socket Direct	81
11.2	Down Stream Port (DSP)	82
11.3	Channel Insertion Loss	83
11.4	Cabline CA-II Plus Harness Pinouts	83
11.4.1	Cabline CA-II Plus Harness - Component Side	83
11.4.2	Cabline CA-II Plus Harness - Print Side	90
11.5	Technical Specifications	100
11.5.1	PCIe Auxiliary Card Mechanical Drawings and Dimensions.....	101
11.5.2	Bracket Mechanical Drawings and Dimensions	101
11.5.3	Cabline CA-II Plus Harnesses Mechanical Drawing	102
12	Supported Servers and Power Cords	103
12.1	Supported Servers	103
12.2	Supported Power Cords	103

13	Document Revision History	104
----	---------------------------------	-----

You can download a PDF version [here](#).

About This Manual

The NVIDIA® BlueField® networking platform ignites unprecedented innovation for modern data centers and supercomputing clusters. With its robust compute power and integrated software-defined hardware accelerators for networking, storage, and security, BlueField creates a secure and accelerated infrastructure for any workload in any environment, ushering in a new era of accelerated computing and AI.

This User Manual describes NVIDIA BlueField-3 DPUs (Data Processing Unit) and SuperNICs. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and a step-by-step plan of how to bring up the BlueField-3 DPUs and SuperNICs.

Ordering Part Numbers

The tables below list the ordering part numbers (OPNs) for available BlueField-3 platforms in Full-Height Half-Length (FHHL) and Half-Height Half-Length (HHHL) form factors.

The Device ID of all DPUs is 41692. All DPUs/SuperNICs are shipped with a tall bracket.

BlueField-3 DPUs

Model and Form Factor	NVIDIA OPN	Series / Cores	Data Transmission Rate	No. of Ports	PCIe Support	x16 PCIe Extension Option	External Power Connector	Crypto	On-Board DDR5 Memory	Integrated BMC	PSID	Life cycle
B3240 Dual-Slot FHHL	900-9D 3B6-00 CN-AB0	P-Series / 16 Arm-Cores	InfiniBand: NDR 400Gb/s (Default) Ethernet : 400GbE	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	✓	32GB	✓	MT_00 00000 883	Mass Production
	900-9D 3B6-00 SN-AB0	P-Series / 16 Arm-Cores	InfiniBand: NDR 400Gb/s (Default) Ethernet : 400GbE	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	-	32GB	✓	MT_00 00000 964	Mass Production
B3240 Single-Slot FHHL For Cold-Aisle Servers Only	900-9D 3B6-00 CN-PA0	P-Series / 16 Arm-Cores	InfiniBand: NDR 400Gb/s (Default) Ethernet : 400GbE	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	✓	32GB	✓	MT_00 00001 188	Prototype

Model and Form Factor	NVIDIA OPN	Series / Cores	Data Transmission Rate	No. of Ports	PCIe Support	x16 PCIe Extension Option	External Power Connector	Crypto	On-Board DDR5 Memory	Integrated BMC	PSID	Lifecycle
B3220 Single-Slot FHHL	900-9D3B6-00 CV-AA0	P-Series / 16 Arm-cores	InfiniBand: NDR200 200Gb/s Ethernet : 200GbE (Default)	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	✓	32GB	✓	MT_000000884	Mass Production
	900-9D3B6-00 SV-AA0	P-Series / 16 Arm-cores	InfiniBand: NDR200 200Gb/s Ethernet : 200GbE (Default)	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	-	32GB	✓	MT_000000965	Mass Production
B3210E Single-Slot FHHL	900-9D3B6-00 CC-EA0	E-Series / 16 Arm-cores	InfiniBand: HDR100 100Gb/s Ethernet : 100GbE (Default)	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	✓	32GB	✓	MT_000000115	Mass Production
	900-9D3B6-00 SC-EA0	E-Series / 16 Arm-cores	InfiniBand: HDR100 100Gb/s Ethernet : 100GbE (Default)	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	✓	-	32GB	✓	MT_000000117	Mass Production

BlueField-3 SuperNICs

Model and Form Factor	NVIDIA OPN	Series/ Cores	Data Transmission Rate	No. of Ports	PCIe Support	Crypto	On-Board DDR5 Memory	Integrated BMC	PSID	Lifecycle
B3210L Single-Slot FHHL	900-9D3B4-00CC-EA0	E-Series / 8 Arm-Cores	InfiniBand: HDR100 100Gb/s Ethernet: 100GbE (Default)	2-Ports QSFP12	PCIe Gen 5.0 x16	✓	16GB	✓	MT_000000966	Mass Production

Model and Form Factor	NVIDIA OPN	Series/ Cores	Data Transmission Rate	No. of Ports	PCIe Support	Crypto	On-Board DDR 5 Memory	Integrated BMC	PSID	Lifecycle
	900-9D3B 4-00SC-EA0	E-Series / 8 Arm-Cores	InfiniBand: HDR100 100Gb/s Ethernet: 100GbE (Default)	2-Ports QSFP112	PCIe Gen 5.0 x16	-	16GB	✓	MT_000 0000967	Mass Production
B3220L Single-Slot FHHL	900-9D3B 4-00CV-EA0	E-Series / 8 Arm-Cores	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE (Default)	2-Ports QSFP112	PCIe Gen 5.0 x16	✓	16GB	✓	MT_000 0001093	Mass Production
	900-9D3B 4-00SV-EA0	E-Series / 8 Arm-Cores	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE (Default)	2-Ports QSFP112	PCIe Gen 5.0 x16	-	16GB	✓	MT_000 0001094	Mass Production
B3140L Single-Slot FHHL	900-9D3B 4-00EN-EA0	E-Series / 8 Arm-Cores	InfiniBand: NDR 400Gb/s (Default) Ethernet: 400GbE	1-Port QSFP112	PCIe Gen 5.0 x16	✓	16GB	✓	MT_000 0001010	Mass Production
	900-9D3B 4-00PN-EA0	E-Series / 8 Arm-Cores	InfiniBand: NDR 400Gb/s (Default) Ethernet: 400GbE	1-Port QSFP112	PCIe Gen 5.0 x16	-	16GB	✓	MT_000 0001011	Mass Production
B3140H Single-Slot HHHL	900-9D3D 4-00EN-HA0	E-Series / 8 Arm-Cores	InfiniBand: NDR 400Gb/s Ethernet: 400GbE (Default)	1-Port QSFP112	PCIe Gen 5.0 x16	✓	16GB	✓	MT_000 0001010	Mass Production
	900-9D3D 4-00NN-HA0	E-Series / 8 Arm-Cores	InfiniBand: NDR 400Gb/s Ethernet: 400GbE (Default)	1-Port QSFP112	PCIe Gen 5.0 x16	-	16GB	✓	MT_000 0001070	Mass Production

EOL'ed (End of Life) DPUs

Intended Audience

This manual is intended for the installer and user of these cards. The manual assumes basic familiarity with InfiniBand/Ethernet network and architecture specifications.

Technical Support

Customers who purchased NVIDIA products directly from NVIDIA are invited to contact us through the following methods:

- URL: www.nvidia.com → Support
- E-mail: enterprisesupport@nvidia.com

Customers who purchased NVIDIA M-1 Global Support Services, please see your contract for details regarding Technical Support.

Customers who purchased NVIDIA products through an NVIDIA-approved reseller should first seek assistance through their reseller.

Related Documentation

InfiniBand Architecture Specification	InfiniBand Trade Association (IBTA) InfiniBand® specification Release 1.3.1, November 2, 2016 and Vol. 2, Release 1.4 , and Vol 2 - Release 1.5.
IEEE Std 802.3 Specification	IEEE Ethernet specification.
PCI Express Specifications	Industry Standard PCI Express Base and Card Electromechanical Specifications.
NVIDIA LinkX Interconnect Solutions	The NVIDIA® LinkX® product family of cables and transceivers provide the industry's broadest portfolio of QDR/FDR10 (40Gb/s), FDR (56Gb/s), EDR/HDR100 (100Gb/s), HDR (200Gb/s) and NDR (400Gb/s) cables, including Direct Attach Copper cables (DACs), copper splitter cables, Active Optical Cables (AOCs) and transceivers in a wide range of lengths from 0.5m to 10km. In addition to meeting IBTA standards, NVIDIA tests every product in an end-to-end environment ensuring a Bit Error Rate of less than 1E-15.
BlueField DPU Platform BSP Documentation	This guide provides product release notes as well as information on the BSP and how to develop and/or customize applications, system software, and file system images for the BlueField platform.
DOCA SDK Software Documentation	NVIDIA DOCA SDK software.

Document Conventions

When discussing memory sizes, GB and GBytes are used in this document to mean size in gigabytes. The use of Gb or Gbits (small b) indicates size in gigabits. In this document PCIe is used to mean PCI Express.

Revision History

A list of the changes made to this document are provided in [Document Revision History](#).

1 Introduction

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.

1.1 System Requirements

Item	Description
PCI Express slot	In PCIe x16 Configuration PCIe Gen 5.0 (32GT/s) through x16 edge connector. In PCIe x16 Extension Option - Switch DSP (Data Stream Port) • PCIe Gen 5.0 SERDES @32GT/s through edge connector • PCIe Gen 5.0 SERDES @32GT/s through PCIe Auxiliary Connection Card
System Power Supply	Minimum 75W or greater system power supply for all cards. B3240, B3220, B3210 and B3210E DPUs require a supplementary 8-pin ATX power supply connectivity through the external power supply connector. NOTE: The power supply harness is not included in the package. Refer to the Hardware Installation and PCIe Bifurcation for important notes and warnings on powering up and down the card.
Operating System	BlueField-3 platforms is shipped with Ubuntu - a Linux commercial operating system - which includes the NVIDIA OFED stack (MLNX_OFED), and is capable of running all customer-based Linux applications seamlessly. For more information, please refer to the DOCA SDK documentation or NVIDIA BlueField DPU BSP.
Connectivity	• Interoperable with 1/10/25/40/50/100/200/400 Gb/s Ethernet switches and SDR/FDR/EDR/HDR100/HDR/NDR200/NDR InfiniBand switches • Passive copper cable with ESD protection • Powered connectors for optical and active cable support

For detailed information, see [Specifications](#).

1.2 Package Contents

Prior to unpacking your product, it is important to make sure your server meets all the system requirements listed above for a smooth installation. Be sure to inspect each piece of equipment shipped in the packing box. If anything is missing or damaged, contact your reseller.

1.2.1 Card Package

For B3240, B3220 and B3210E DPUs, you need an 8-pin PCIe external power cable to activate the card. The cable is not included in the package. For further details, please refer to [External PCIe Power Supply Connector](#).

Item	Description
Card	1x BlueField-3 platform
Accessories	1x tall bracket (shipped assembled on the card)

1.2.2 Accessories Kit

This is an optional accessories kit used for debugging purposes and can be ordered separately.

Kit OPN	Contents
MBF35-DKIT	4-pin USB to female USB Type-A cable
	20-pin shrouded connector to USB Type-A cable

1.2.3 Optional PCIe Auxiliary Card Package

The Socket-Direct functionality is currently not supported by firmware. Please approach your sales representatives.

This is an optional kit which applies to following OPNs:

- B3220: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
- B3240: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0
- B3210: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- B3210E: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

The PCIe auxiliary kit is purchased separately to utilize the Socket-Direct functionality in dual-socket servers or for downstream port extension option. For package contents and more information, refer to [PCIe Auxiliary Card Kit](#).

1.3 Features and Benefits

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

Feature	Description				
InfiniBand Architecture Specification v1.5 compliant	BlueField-3 platforms deliver low latency, high bandwidth, and computing efficiency for high-performance computing (HPC), artificial intelligence (AI), and hyperscale cloud data centers applications.				
	BlueField-3 platforms are InfiniBand Architecture Specification v1.5 compliant.				
	InfiniBand Network Protocols and Rates:				
	Protocol	Standard	Rate (Gb/s)		Comments
			4x Port (4 Lanes)	2x Ports (2 Lanes)	
	NDR/NDR200	IBTA Vol2 1.5	425	212.5	PAM4 256b/257b encoding and RS-FEC
	HDR/HDR100	IBTA Vol2 1.4	212.5	106.25	PAM4 256b/257b encoding and RS-FEC
EDR	IBTA Vol2 1.3.1	103.125	51.5625	NRZ 64b/66b encoding	
FDR	IBTA Vol2 1.2	56.25	N/A	NRZ 64b/66b encoding	

Feature	Description																								
Up to 400 Gigabit Ethernet	BlueField-3 platforms comply with the following IEEE 802.3 standards: 400GbE / 200GbE / 100GbE / 50GbE / 40GbE / 25GbE / 10GbE																								
	<table><tr><th>Protocol</th><th>MAC Rate</th></tr><tr><td>IEEE802.3ck</td><td>400/200/100 Gigabit Ethernet (Include ETC enhancement)</td></tr><tr><td>IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu</td><td>400/200/100 Gigabit Ethernet (Include ETC enhancement)</td></tr><tr><td>IEEE 802.3bj IEEE 802.3bm</td><td>100 Gigabit Ethernet</td></tr><tr><td>IEEE 802.3by Ethernet Consortium25</td><td>50/25 Gigabit Ethernet</td></tr><tr><td>IEEE 802.3ba</td><td>40 Gigabit Ethernet</td></tr><tr><td>IEEE 802.3ae</td><td>10 Gigabit Ethernet</td></tr><tr><td>IEEE 802.3cb</td><td>2.5/5 Gigabit Ethernet (For 2.5: support only 2.5 x1000BASE-X)</td></tr><tr><td>IEEE 802.3ap</td><td>Based on auto-negotiation and KR startup</td></tr><tr><td>IEEE 802.3ad IEEE 802.1AX</td><td>Link Aggregation</td></tr><tr><td>IEEE 802.1Q IEEE 802.1P VLAN tags and priority</td><td></td></tr><tr><td>IEEE 802.1Qau (QCN) Congestion Notification IEEE 802.1Qaz (ETS) EEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)</td><td></td></tr></table>	Protocol	MAC Rate	IEEE802.3ck	400/200/100 Gigabit Ethernet (Include ETC enhancement)	IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu	400/200/100 Gigabit Ethernet (Include ETC enhancement)	IEEE 802.3bj IEEE 802.3bm	100 Gigabit Ethernet	IEEE 802.3by Ethernet Consortium25	50/25 Gigabit Ethernet	IEEE 802.3ba	40 Gigabit Ethernet	IEEE 802.3ae	10 Gigabit Ethernet	IEEE 802.3cb	2.5/5 Gigabit Ethernet (For 2.5: support only 2.5 x1000BASE-X)	IEEE 802.3ap	Based on auto-negotiation and KR startup	IEEE 802.3ad IEEE 802.1AX	Link Aggregation	IEEE 802.1Q IEEE 802.1P VLAN tags and priority		IEEE 802.1Qau (QCN) Congestion Notification IEEE 802.1Qaz (ETS) EEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)	
	Protocol	MAC Rate																							
	IEEE802.3ck	400/200/100 Gigabit Ethernet (Include ETC enhancement)																							
	IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu	400/200/100 Gigabit Ethernet (Include ETC enhancement)																							
	IEEE 802.3bj IEEE 802.3bm	100 Gigabit Ethernet																							
	IEEE 802.3by Ethernet Consortium25	50/25 Gigabit Ethernet																							
	IEEE 802.3ba	40 Gigabit Ethernet																							
	IEEE 802.3ae	10 Gigabit Ethernet																							
	IEEE 802.3cb	2.5/5 Gigabit Ethernet (For 2.5: support only 2.5 x1000BASE-X)																							
	IEEE 802.3ap	Based on auto-negotiation and KR startup																							
	IEEE 802.3ad IEEE 802.1AX	Link Aggregation																							
	IEEE 802.1Q IEEE 802.1P VLAN tags and priority																								
	IEEE 802.1Qau (QCN) Congestion Notification IEEE 802.1Qaz (ETS) EEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)																								
	On-board Memory	• Quad SPI NOR FLASH - includes 256Mbit for Firmware image. • UVPS EEPROM - includes 2Mbit. • FRU EEPROM - Stores the parameters and personality of the card. The EEPROM capacity is 128Kbit. FRU I2C address is (0x50) and is accessible through the PCIe SMBus. • DPU_BMC Flashes: • 2x 64MByte for BMC Image • 512MByte for Config Data • eMMC pSLC 40GB with 30K Write Cycles eMMC for SoC BIOS. • SSD (onboard BGA) 128GByte for user SoC OS, logs and application SW. • DDR5 SDRAM - 16GB/32GB @5200MT/s or @5600MT/s single/dual-channel DDR5 SDRAM memory. Solder down on-board. 128bit + 16bit ECC																							
BlueField-3 IC	The BlueField-3 platforms integrate x8 / x16 Armv8.2+ A78 Hercules cores (64-bit) is interconnected by a coherent mesh network, one DRAM controller, an RDMA intelligent network adapter supporting up to 400Gb/s, an embedded PCIe switch with endpoint and root complex functionality, and up to 32 lanes of PCIe Gen 5.0.																								

Feature	Description
Overlay Networks	In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. BlueField-3 platforms effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.
RDMA and RDMA over Converged InfiniBand/Ethernet (RoCE)	Utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged InfiniBand/Ethernet) technology, the BlueField-3 platforms deliver low-latency and high-performance over InfiniBand/Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.
Quality of Service (QoS)	Support for port-based Quality of Service enabling various application requirements for latency and SLA.
Storage Acceleration	- A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage RDMA for high-performance storage access: NVMe over Fabric offloads for the target machine - BlueField-3 cards 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 the NVIDIA BlueField-3 Arm cores. On the storage initiator side, BlueField-3 networking platforms can prove an efficient solution for hyper-converged systems to enable the host CPU to focus on compute while all the storage interface is handled through the Arm cores.
NVMe-oF	Non-volatile Memory Express (NVMe) over Fabrics is a protocol for communicating block storage IO requests over RDMA to transfer data between a host computer and a target solid-state storage device or system over a network. BlueField-3 platforms may operate as a co-processor offloading specific storage tasks from the host using its powerful NVMe over Fabrics Offload accelerator.
SR-IOV	The SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
High-Performance Accelerations	- Adaptive Routing on Reliable Transport - Burst Buffer Offloads for Background Checkpointing
GPU Direct	GPUDirect RDMA is a technology that provides a direct P2P (Peer-to-Peer) data path between the GPU Memory directly to/from the NVIDIA HCA devices. This provides a significant decrease in GPU-GPU communication latency and completely offloads the CPU, removing it from all GPU-GPU communications across the network. BlueField-3 platforms use high-speed DMA transfers to copy data between P2P devices resulting in more efficient system applications
Isolation	BlueField-3 platforms function 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-3 may detect/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-3 hardware-based accelerations offload the crypto operations and free up the CPU, reducing latency and enabling scalable crypto solutions. BlueField-3 “host-unaware” solutions may transmit and receive data, while BlueField-3 acts as a bump-in-the-wire for crypto.

Feature	Description
Security Accelerators	A consolidated compute and network solution based on BlueField-3 achieves significant advantages over a centralized security server solution. Standard encryption protocols and security applications can leverage BlueField-3 compute capabilities and network offloads for security application solutions such as Layer4 Statefull Firewall.
Virtualized Cloud	By leveraging BlueField-3 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
Out-of-Band Management	The BlueField-3 platforms incorporate a 1GbE RJ45 out-of-band port that allows the network operator to establish trust boundaries in accessing the management function to apply it to network resources. It can also be used to ensure management connectivity (including the ability to determine the status of any network component) independent of the status of other in-band network components.
BMC	Some BlueField-3 platforms incorporate local NIC BMC (Baseboard Management Controller) hardware on the board. The BMC SoC (system on a chip) can utilize either shared or dedicated NICs for remote access. The BMC node enables remote power cycling, board environment monitoring, BlueField-3 chip temperature monitoring, board power and consumption monitoring, and individual interface resets. The BMC also supports the ability to push a bootstream to BlueField-3. Having a trusted on-board BMC that is fully isolated for the host server ensures highest security for the BlueField-3 platforms.

2 BlueField-3 Administrator Quick Start Guide

This page is tailored for system administrators wishing to install BlueField and perform sample administrative actions on it. For a quick start guide aimed at software developers wishing to develop applications on the BlueField card using the DOCA framework, please refer to the [NVIDIA DOCA Developer Quick Start Guide](#).

Not sure which guide to follow? For more details on the different BlueField user types, please refer to the [NVIDIA BlueField and DOCA User Types](#) document.

2.1 Prerequisites for Initial BlueField-3 Deployment

Unable to render include or excerpt-include. Could not retrieve page.

2.2 First-time Installation Procedure



Unable to render include or excerpt-include. Could not retrieve page.

3 Supported Interfaces

This section describes the supported interfaces. The table below describes each numbered interface referenced in the figures, with a link to detailed information.





The below figures are for illustration purposes only and might not reflect the current revision of the BlueField-3 platforms.

3.1 BlueField-3 SuperNICs Layout and Interface Information

OPN	SuperNIC Component Side	SuperNIC Print Side
HHHL Single-Slot SuperNIC Model: B3140H 900-9D3D4-00EN-HA0 900-9D3D4-00NN-HA0	 A top-down view of the SuperNIC component side of the B3140H model. It shows various components labeled with numbers 1 through 12. Labels 1-4 point to connectors on the bottom edge, 5 points to the central SoC, 6 points to a USB connector, 7 points to a management interface, 8 points to a 1GbE OOB interface, 9 points to the integrated BMC, 10 points to the SSD, 11 points to the PCIe connector, and 12 points to the network ports.	 A top-down view of the SuperNIC print side of the B3140H model. It shows the component layout from the opposite side, with label 5 pointing to the central SoC.

Item	Interface	Description
1	SoC	8 Arm-Cores SuperNIC SoC (System-on-Chip)
2	Networking Interface	The network traffic is transmitted through the SuperNIC QSFP112 connectors. The QSFP112 connectors allow the use of modules and optical and passive cable interconnect solutions
3	Networking Ports LEDs Interface	One bi-color I/O LEDs per port to indicate link and physical status
4	PCI Express Interface	PCIe Gen 5.0 through an x16 edge connector
5	DDR5 SDRAM On-Board Memory	Single-Channel Cards: 10 units of DDR5 SDRAM for a total of 16GB @ 5200MT/s 64bit + 8bit ECC, solder-down memory
6	NC-SI Management Interface	NC-SI 20 pins BMC connectivity for remote management
7	USB 4-pin RA Connector	Used for OS image loading
8	1GbE OOB Management Interface	1GbE BASE-T OOB management interface
9	Integrated BMC	SuperNIC BMC
10	SSD Interface	128GB







Item	Interface	Description
11	RTC Battery	Battery holder for RTC
12	eMMC	x8 NAND flash

OPN	SuperNIC Component Side	SuperNIC Print Side
FHHL Single-Slot Dual-Port SuperNICs Model: B3220L 900-9D3B4-00CV-EA0 900-9D3B4-00SV-EA0 Model: B3210L 900-9D3B4-00CC-EA0 900-9D3B4-00SC-EA0		
FHHL Single-Slot Single-Port SuperNICs Model: B3140L 900-9D3B4-00EN-EA0 900-9D3B4-00PN-EA0		

Item	Interface	Description
1	SoC	8/16 Arm-Cores SuperNIC SoC
2	Networking Interface	The network traffic is transmitted through the QSFP112 connectors. The QSFP112 connectors allow the use of modules and optical and passive cable interconnect solutions
3	Networking Ports LEDs Interface	One bi-color I/O LEDs per port to indicate link and physical status
4	PCI Express Interface	PCIe Gen 5.0/4.0 through an x16 edge connector
5	DDR5 SDRAM On-Board Memory	20 units of DDR5 SDRAM for a total of 32GB @5200 or 5600MT/s. 128bit + 16bit ECC, solder-down memory
6	NC-SI Management Interface	NC-SI 20 pins BMC connectivity for remote management
7	USB 4-pin RA Connector	Used for OS image loading
8	1GbE OOB Management Interface	1GbE BASE-T OOB management interface

Item	Interface	Description
9	MMCX RA PPS IN/OUT	Allows PPS IN/OUT
12	Integrated BMC	SuperNIC BMC
13	SSD Interface	128GB
14	RTC Battery	Battery holder for RTC
15	eMMC	x8 NAND flash

3.2 BlueField-3 DPUs Layout and Interface Information

OPN	DPU Component Side	DPU Print Side
FHHL Single-Slot Dual-Port DPUs with PCIe Extension Option Model: B3210E 900-9D3B6-00CC-EA0 900-9D3B6-00SC-EA0 Model: B3210 900-9D3B6-00CC-AA0 900-9D3B6-00SC-AA0 Model: B3220 900-9D3B6-00CV-AA0 900-9D3B6-00SV-AA0		
FHHL Dual-Slot Dual-Port DPUs Model: B3240 900-9D3B6-00CN-AB0 900-9D3B6-00SN-AB0		
FHHL Cold-Aisle Single-Slot Dual-Port DPUs Model: B3240 900-9D3B6-00CN-PA0 This product can operate in cold-aisle servers only.		

Item	Interface	Description
1	DPU SoC	8/16 Arm-Cores DPU SoC
2	Networking Interface	The network traffic is transmitted through the DPU QSFP112 connectors. The QSFP112 connectors allow the use of modules and optical and passive cable interconnect solutions

Item	Interface	Description
3	Networking Ports LEDs Interface	One bi-color I/O LEDs per port to indicate link and physical status
4	PCI Express Interface	PCIe Gen 5.0/4.0 through an x16 edge connector
5	DDR5 SDRAM On-Board Memory	20 units of DDR5 SDRAM for a total of 32GB @5200 or 5600MT/s. 128bit + 16bit ECC, solder-down memory
6	NC-SI Management Interface	NC-SI 20 pins BMC connectivity for remote management
7	USB 4-pin RA Connector	Used for OS image loading
8	1GbE OOB Management Interface	1GbE BASE-T OOB management interface
9	MMCX RA PPS IN/OUT	Allows PPS IN/OUT
10	External PCIe Power Supply Connector	An external 12V power connection through an 8-pin ATX connector Applies to models: B3210E, B3210 and B3220
11	Cabline CA-II Plus Connectors	Two Cabline CA-II plus connectors are populated to allow connectivity to an additional PCIe x16 Auxiliary card Applies to models: B3210E, B3210 and B3220
12	Integrated BMC	DPU BMC
13	SSD Interface	128GB
14	RTC Battery	Battery holder for RTC
15	eMMC	x8 NAND flash

3.3 Interfaces Detailed Description

3.3.1 System-on-Chip (SoC)

NVIDIA® BlueField®-3 is a family of advanced IC solutions that integrate a coherent mesh of 64-bit Armv8.2+ A78 Hercules cores, an NVIDIA® ConnectX®-7 network adapter front-end, and a PCI Express switch into a single chip. The powerful SoC architecture includes an Armv multicore processor array, enabling customers to develop sophisticated applications and highly differentiated feature sets. Leverages the rich Arm software ecosystem and introduces the ability to offload the x86 software stack.

At the heart of BlueField-3, the ConnectX-7 network offload controller with RDMA and RDMA over Converged Ethernet (RoCE) technology 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 encaps/decaps of NVGRE, VXLAN, and MPLS overlay protocols.

3.3.1.1 Encryption

Applies to Crypto enabled OPNs.


The BlueField-3 networking platforms address the concerns of modern data centers by combining hardware encryption accelerators with embedded software and fully integrated advanced network capabilities, making it ideal platforms for developing proprietary security applications. The platforms enable a distributed security architecture by isolating and protecting each workload and provide flexible control and visibility at the server and workload level; controlling risk at the server access layer builds security into the DNA of the data center and enables prevention, detection, and response to potential threats in real-time. The BlueField-3 platforms can deliver powerful functionality, including encryption of data-in-motion, bare-metal provisioning, stateful L4 firewall, and more.

3.3.2 Networking Interface

The networking platforms include special circuits to protect the card/server from ESD shocks when plugging copper cables.

The network ports are compliant with the InfiniBand Architecture Specification, Release 1.5. InfiniBand traffic is transmitted through the cards' QSFP112 connectors.

3.3.3 Networking Ports LEDs Interface



One bicolor (Yellow and Green) I/O LED per port indicates speed and link status.

Link Indications

State	Bi-Color LED (Yellow/Green)	
Beacon command for locating the adapter card	1Hz blinking Yellow	

State	Bi-Color LED (Yellow/Green)										
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following: <table> <tr> <th>Error Type</th><th>Description</th><th>LED Behavior</th></tr> <tr> <td>I²C</td><td>I²C access to the networking ports fails</td><td>Blinks until error is fixed</td></tr> <tr> <td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until error is fixed</td></tr> </table>	Error Type	Description	LED Behavior	I ² C	I ² C access to the networking ports fails	Blinks until error is fixed	Over-current	Over-current condition of the networking ports	Blinks until error is fixed	
Error Type	Description	LED Behavior									
I ² C	I ² C access to the networking ports fails	Blinks until error is fixed									
Over-current	Over-current condition of the networking ports	Blinks until error is fixed									
Physical Activity	Blinking Green										
Link Up	Solid Green										
Physical Up (InfiniBand Mode Only)	Solid Yellow										

3.3.4 PCI Express Interface

The BlueField-3 supports PCI Express Gen 5.0/4.0 through x16 edge connectors. Some cards allow connectivity to an additional PCIe x16 Auxiliary card through the Cabline CA-II Plus connectors.

The following lists PCIe interface features:

- PCIe Gen 5.0, 4.0, 3.0, 2.0 and 1.1 compatible
- 2.5, 5.0, or 8.0, 16.0 or 32.0 GT/s link rate x16 lanes
- Auto-negotiates to x16, x8, x4, x2, or x1

3.3.5 DDR5 SDRAM On-Board Memory

The BlueField-3 incorporate 10 or 20 units of DDR5 SDRAM. See the following table for DDR5 SDRAM memory specifications per ordering part number.

Model	OPNs	DDR5 SDRAM On-Board Memory
B3140L B3140H B3210L B3220L	900-9D3B4-00EN-EA0 / 900-9D3B4-00PN-EA0 900-9D3D4-00EN-HA0 / 900-9D3D4-00NN-HA0 900-9D3B4-00CC-EA0 / 900-9D3B4-00SC-EA0 900-9D3B4-00CV-EA0 / 900-9D3B4-00SV-EA0	Single-channel with 10 DDR5 + ECC (64bit + 8bit ECC) for a total of 16GB @ 5200MT/s

Model	OPNs	DDR5 SDRAM On-Board Memory
B3210E B3220^(a)	900-9D3B6-00CC-EA0 / 900-9D3B6-00SC-EA0 900-9D3B6-00CV-AA0 / 900-9D3B6-00SV-AA0	Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5200MT/s
B3240	900-9D3B6-00CN-AB0 / 900-9D3B6-00SN-AB0 / 900-9D3B6-00CN-PA0	Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5600MT/s

a. 5600MT/s is supported from HW rev (AC). For more information, refer to the latest [NVIDIA BlueField DPU BSP](#).

3.3.6 NC-SI Management Interface

BlueField-3 enables the connection of a Baseboard Management Controller (BMC) to a set of Network Interface Controller (NICs) to enable out-of-band remote manageability. The NC-SI management is supported over RMII and has a connector on the card. Please refer to [NC-SI Management Interface](#) for pins.

The below table specifies the maximum trace lengths per board type. Please take the maximum trace length on the board into consideration in your design.

The USB to UART cable is not used for NC-SI management purposes.

BlueField-3 Platform Family	SKUs	Maximum Trace Length on the Board
B3140H	900-9D3D4-00EN-HA0 / 900-9D3D4-00NN-HA0	82.063mm (3.2inch)
B3220	900-9D3B6-00CV-AA0 / 900-9D3B6-00SV-AA0	144.449mm (~5.7inch)
B3210	900-9D3B6-00CC-AA0 / 900-9D3B6-00SC-AA0	
B3240	900-9D3B6-00CN-AB0 / 900-9D3B6-00SN-AB0 / 900-9D3B6-00CN-PA0	
B3210E	900-9D3B6-00CC-EA0 / 900-9D3B6-00SC-EA0	
B3210L	900-9D3B4-00CC-EA0 / 900-9D3B4-00SC-EA0	
B3220L	900-9D3B4-00CV-EA0 / 900-9D3B4-00SV-EA0	
B3140L	900-9D3B4-00EN-EA0 / 900-9D3B4-00PN-EA0	

3.3.7 UART Interface Connectivity

The UART debug interface on BlueField-3 boards can be accessed through a 20-pin NC-SI connector, which is associated with the NIC BMC device. The connectivity is shown in the following table:

NC-SI Connector Pin #	Signal on Board
12	GND
14	UART_TX
16	UART_RX

It is prohibited to connect any RS-232 cable directly! Only TTL 3.3V voltage level cables are supported.

Do not use the USB-to-UART cable for NC-SI management purposes.

3.3.8 USB 4-pin RA Connector

The USB 4-pin RA USB connector is used to load operating system images. Use a 4-pin male connector to a male Type-A cable to connect to the board.



It is prohibited to connect male-to-male to host, it is only used for a disk on key.

The male connector to the male Type-A cable is not included in the shipped card box and should be ordered separately as part of the accessories kit (P/N: MBF35-DKIT).

3.3.9 1GbE OOB Management Interface

BlueField-3 incorporates a 1GbE RJ45 out-of-band port that allows the network operator to establish trust boundaries in accessing the management function to apply it to network resources. It can also be used to ensure management connectivity (including the ability to determine the status of any network component) independent of other in-band network components' status.

For cards with integrated BMC: 1GbE OOB Management can be performed via the integrated BMC.

3.3.9.1 1GbE OOB Management LEDs Interface

Two OOB management LEDs, one Green and one Yellow, behave as described in the table below.

Green LED	Yellow LED	Link/Activity
OFF	OFF	Link off
ON	OFF	1 Gb/s link / No activity
Blinking	OFF	1 Gb/s link / Activity (RX,TX)
OFF	ON	Not supported
OFF	Blinking	
ON	ON	
Blinking	Blinking	

3.3.10 PPS IN/OUT Interface

BlueField-3 incorporates an integrated Hardware Clock (PHC) that allows the card to achieve sub-20u Sec accuracy and also offers many timing-related functions such as time-triggered scheduling or time-based SND accelerations (time-based ASAP²). Furthermore, 5T technology enables the software application to transmit fronthaul (ORAN) at high bandwidth. The PTP part supports the subordinate clock, master clock, and boundary clock.

BlueField-3 PTP solution allows you to run any PTP stack on your host.

With respect to testing and measurements, selected BlueField-3 boards allow you to use the PPS-out signal from the onboard MMCX RA connector. The BlueField-3 board also allow measuring PTP in scale with the PPS-In signal. The PTP HW clock on the Network adapter is sampled on each PPS-In signal, and the timestamp is sent to the SW.

3.3.11 External PCIe Power Supply Connector

Applies to following DPUs only. The external ATX power cable is not supplied with the DPU package; however, this is a standard cable usually available in servers.

- B3220: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
- B3240: 900-9D3B6-00CN-AB0, 900-9D3B6-00SN-AB0 and 900-9D3B6-00CN-PA0
- B3210: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- B3210E: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

To power up the above-mentioned DPUs, it is necessary to use a supplementary 8-pin ATX power cable. Since the power provided by the PCIe golden fingers is limited to 66W, a total maximum of up to 150W is enabled through the ATX 8-pin connector and the PCIe x16 golden fingers.


The maximum power consumption which does not exceed 150W, is in accordance with the mode of operation of the DPU, and is split between the two power sources as follows:

- Up to 66W from the PCIe golden fingers (12V)
- The rest of the consumed power is drawn from the external PCIe power supply connector

Important Notes and Warnings

- The BlueField-3 DPU requires a standard PCIe ATX power connection. However, certain servers may require a custom setup to enable ATX power compatibility.
- Consult the manufacturer's manuals and specifications to ensure proper connectivity.
- Before connecting the ATX power cable, make sure you are using a PCIe-compliant 8-pin ATX cable from the server power supply kit.
- Do not link the CPU power cable to the BlueField-3 DPU PCIe ATX power connector, as their pin configurations differ. Using the CPU power cable in this manner is strictly prohibited and can potentially damage the BlueField-3 DPU. Please refer to [External PCIe Power Supply Connector Pins](#) for the external PCIe power supply pins.
- Avoid using non-standard cables that do not comply with the DPU, unnecessary adapter cables, or storing the cables near heat sources.
- It is preferable that the x16 PCIe golden fingers and the PCI ATX power supply draw from the same power source. For more information on how to power up the card, refer to [DPU Power-Up Instructions](#).
- The PCIe ATX 8-pin connector is not compatible with an EPS12V power cable source. Ensure that the appropriate PCIe auxiliary power source is available, not an EPS12V power source.
- If you are uncertain about your server's compatibility with the PCI ATX connection, please contact your NVIDIA representative for assistance.

Mechanical Keys of the 8-pin ATX Power Cable (PCIe CEM Specification Rev 5.0)



3.3.12 Cabline CA-II Plus Connectors

Applies to the following DPUs:

B3220: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0

B3240: 900-9D3B6-00CN-AB0, 900-9D3B6-00SN-AB0 and [900-9D3B6-00CN-PA0](#)

B3210: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0

B3210E: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

The Cabline CA-II connectors on the DPU enable connectivity to an additional PCIe x16 bus in addition to the PCIe x16 bus available through the golden-fingers. The Cabline CA-II Plus connectors allow connectivity to flash cards and NVMe SSD drives.

Specific applications have an interest in direct connectivity to the far end of the Cabline CA-II cables, through the two 60-pin Cabline CA-II connectors, directly to the motherboard, in order to cut the insertion loss and/or the additional space associated with a PCIe x16 Flash Auxiliary Board.

The Cabline CA-II connectors mate with two 60-pin Cabline CA-II cables that can be distinguished by their black or white external insulators and connector pinouts. The black Cabline CA-II cable mates with the DPU's component (top) side, whereas the white Cabline CA-II cable mates with the DPU print (bottom) side. The Cabline CA-II cables are offered in three standard lengths; 150mm, 350mm, and 550mm.

For connector pinouts, please refer to [Cabline CA-II Plus Connectors Pinouts](#).

3.3.13 Integrated BMC Interface

The card incorporates an onboard integrated NIC BMC and an Ethernet switch. The BMC becomes available once the host server powers up the card. The NIC BMC can control the DPU's power and enables DPU shutdown and power-up.

3.3.14 NVMe SSD Interface

The Self Encrypting Disk (SED) capability is not supported.

It is important to note that all SSD devices come with a limitation on the total number of write operations they can handle throughout their lifespan. This limit is influenced significantly by the software use case and specific parameters like block size and the pattern of data access (whether it is sequential or random). It is the customer's responsibility to oversee the rate at which the SSD ages during both the validation of the code and its usage in the field, ensuring that it aligns with the intended use case.

The on-board 128GB client-grade NVMe SSD is utilized for persistent storage of user applications and logs.

3.3.15 RTC Battery

BlueField-3 incorporates a coin type Lithium battery CR621 for RTC (Real Time Clock).

3.3.16 eMMC Interface

It is important to note that all eMMC devices come with a limitation on the total number of write operations they can handle throughout their lifespan. This limit is influenced significantly by the software use case and specific parameters like block size and the pattern of data access (whether it is sequential or random). It is the customer's responsibility to oversee the rate at which the eMMC ages during both the validation of the code and its usage in the field, ensuring that it aligns with the intended use case.

The eMMC is an x8 NAND flash used for Arm boot and operating system storage. Memory size is 128GB, where it is effectively pSLC 40GB.

4 Pinouts Description



4.1 PCI Express Interface

The following table lists the PCI Express pins description. For further details, please refer to [PCI Express Interface](#).

DPU PCI Express x16 Pin Description

Pin #	Signal Name	Description	Pin #	Signal Name	Description
A1	PRSNT1#	Mechanical Present	B1	12V	
A2	12V		B2	12V	
A3	12V		B3	12V	
A4	GND		B4	GND	
A5	TCK	JTAG - Not Connected	B5	SMCLK	Host SMBus
A6	TDI	JTAG - Not Connected	B6	SMDAT	Host SMBus
A7	TDO	JTAG - Not Connected	B7	GND	
A8	TMS	JTAG - Not Connected	B8	3.3V	3.3V - (Connected in 900-9D3B4-00CC-EA0 & 900-9D3B4-00SC-EA0 only)
A9	3.3V	3.3V - (Connected in 900-9D3B4-00CC-EA0 & 900-9D3B4-00SC-EA0 only)	B9	TRST#	JTAG - Not Connected
A10	3.3V	3.3V - (Connected in 900-9D3B4-00CC-EA0 & 900-9D3B4-00SC-EA0 only)	B10	3.3V_AUX	
A11	PERST#	PCIe Reset	B11	WAKW#/RSVD	
A12	GND		B12	RSVD	
A13	REFCLK+	Host Reference Clock	B13	GND	
A14	REFCLK-	Host Reference Clock	B14	PETP0	
A15	GND		B15	PETN0	
A16	PERP0		B16	GND	
A17	PERN0		B17	RSVD	
A18	GND		B18	GND	
A19	RSVD		B19	PETP1	
A20	GND		B20	PETN1	
A21	PERP1		B21	GND	
A22	PERN1		B22	GND	
A23	GND		B23	PETP2	
A24	GND		B24	PETN2	
A25	PERP2		B25	GND	

Pin #	Signal Name	Description	Pin #	Signal Name	Description
A26	PERN2		B26	GND	
A27	GND		B27	PETP3	
A28	GND		B28	PETN3	
A29	PERP3		B29	GND	
A30	PERN3		B30	RSVD	
A31	GND		B31	RSVD	
A32	RSVD		B32	GND	
A33	RSVD		B33	PETP4	
A34	GND		B34	PETN4	
A35	PERP4		B35	GND	
A36	PERN4		B36	GND	
A37	GND		B37	PETP5	
A38	GND		B38	PETN5	
A39	PERP5		B39	GND	
A40	PERN5		B40	GND	
A41	GND		B41	PETP6	
A42	GND		B42	PETN6	
A43	PERP6		B43	GND	
A44	PERN6		B44	GND	
A45	GND		B45	PETP7	
A46	GND		B46	PETN7	
A47	PERP7		B47	GND	
A48	PERN7		B48	RSVD	
A49	GND		B49	GND	
A50	RSVD		B50	PETP8	
A51	GND		B51	PETN8	
A52	PERP8		B52	GND	
A53	PERN8		B53	GND	
A54	GND		B54	PETP9	
A55	GND		B55	PETN9	
A56	PERP9		B56	GND	
A57	PERN9		B57	GND	
A58	GND		B58	PETP10	
A59	GND		B59	PETN10	
A60	PERP10		B60	GND	
A61	PERN10		B61	GND	

Pin #	Signal Name	Description	Pin #	Signal Name	Description
A62	GND		B62	PETP11	
A63	GND		B63	PETN11	
A64	PERP11		B64	GND	
A65	PERN11		B65	GND	
A66	GND		B66	PETP12	
A67	GND		B67	PETN12	
A68	PERP12		B68	GND	
A69	PERN12		B69	GND	
A70	GND		B70	PETP13	
A71	GND		B71	PETN13	
A72	PERP13		B72	GND	
A73	PERN13		B73	GND	
A74	GND		B74	PETP14	
A75	GND		B75	PETN14	
A76	PERP14		B76	GND	
A77	PERN14		B77	GND	
A78	GND		B78	PETP15	
A79	GND		B79	PETN15	
A80	PERP15		B80	GND	
A81	PERN15		B81	PRSNT2#	Mechanical Present
A82	GND		B82	GND	

4.2 External Power Supply Connector


The following table provides the External Power Supply pins of the external power supply interfaces on the DPU. For further details, please refer to [External PCIe Power Supply Connector](#).

The mechanical pinout of the 8-pin external +12V power connector is shown below. The +12V connector is a GPU power PCIe standard connector. Care should be taken to ensure the power is applied to the correct pins as some 8-pin ATX-type connectors can have different pinouts.

Pin Number	Description
1	12V
2	12V
3	12V
4	Sense1

Pin Number	Description
5	GND
6	Sense0
7	GND
8	GND

Wiring Instructions



4.3 NC-SI Management Interface

The following table lists the NC-SI management interface pinout descriptions. For further details, please refer to [NC-SI Management Interface](#).

Pin #	Signal Name	I/O	Signal Description
1	GND	GND	Ground
2	PKG_ID1	Input (to BlueField-3)	NC-SI PKG_ID Should be connected to the Primary controller NC-SI PKG_ID pins to set the appropriate package ID. PKG_ID0 should be connected to the endpoint device GPIO associated with Package ID[0]. PKG_ID1 should be associated with Package ID[1]. Baseboard should connect to GND or leave floating. DPU should have a 4.7k PU.
3	RBT_RXD0	Output (from BlueField-3)	Receive data. Data signals from the network controller to the BMC. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor to GND on the baseboard between the BMC and the RBT isolator to prevent the signal from floating when no card is installed. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.

Pin #	Signal Name	I/O	Signal Description
4	RBT_REF_CLK	Input	RBT Reference clock. Synchronous clock reference for receive, transmit and control interface. The clock should have a typical frequency of 50MHz $\pm$50 ppm. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the DPU cable connector. The RBT_REF_CLK should not be driven until 3.3V AUX is present on the DPU. The RBT_REF_CLK should be continuous once it has started. For DPUs, this pin should be connected between the connector and the RBT PHY. No external termination is required.
5	RBT_RXD1	Output	Receive data. Data signals from the network controller to the BMC. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor to GND on the baseboard between the BMC and the RBT isolator to prevent the signal from floating when no card is installed. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.
6	GND	GND	Ground
7	RBT_CRSDV	Output	Carrier sense/receive data valid. This signal is used to indicate to the baseboard that the carrier sense/receive data is valid. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor on the baseboard between the BMC and the RBT isolator to prevent the signal from floating when no DPU is installed. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.
8	RBT_ISOLATE_N	Output	This signal is used to indicate the DPU has powered and is ready for NC-SI physical layer connection to be present. When low the baseboard circuitry will isolate the NC-SI connection to the DPU. When high normal NC-SI RBT connectivity is available. Baseboards should terminate this with a 47K-100K PD resistor. DPUs should terminate with a 10k PU resistor.
9	GND	GND	Ground
10	PKG_ID0	Input	NC-SI PKG_ID should be connected to the Primary controller NC-SI PKG_ID pins to set the appropriate package ID. PKG_ID0 should be connected to the endpoint device GPIO associated with Package ID[0]. PKG_ID1 should be associated with Package ID[1]. Baseboard should connect to GND or leave floating. DPU should have a 4.7k PU.
11	RBT_TX_EN	Input	Transmit enable. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor to ground on the baseboard between the RBT isolator and the DPU cable connector to prevent the card-side signals from floating when the RBT signals are isolated. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.
12	GND	GND	Ground

Pin #	Signal Name	I/O	Signal Description
13	RBT_TXD0	Input	Transmit data. Data signals from the BMC to the network controller. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor to GND on the baseboard between the RBT isolator and the DPU cable connector to prevent the card-side signals from floating when the RBT signals are isolated. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.
14	UART_TX	Input	3.3V UART TX signal from the baseboard
15	RBT_TXD1	Input	Transmit data. Data signals from the BMC to the network controller. For baseboards, this pin should be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kΩ pull down resistor to GND on the baseboard between the RBT isolator and the DPU cable connector to prevent the card-side signals from floating when the RBT signals are isolated. For DPUs, this pin should be connected between the connector and the RBT PHY. External termination determined by the DPU RBT PHY requirements.
16	UART_RX	Output	3.3V UART RX signal to the baseboard
17	PRESENCE_N		Presence of DPU. Baseboard should implement a 200 Ω series resistor and 4.7kohm pull-up resistor to 3.3V AUX. DPU should tie this to GND.
18	GND	GND	Ground
19	RBT_ARB_OUT	Input	NC-SI hardware arbitration output. If the baseboard supports multiple DPUs cards connected to the same RBT interface, it should implement logic that connects the RBT_ARB_OUT pin of the first populated DPU card to its RBT_ARB_IN pin if it is the only card present or to the RBT_ARB_IN pin of the next populated card and so on sequentially for all cards on the specified RBT bus to ensure the arbitration ring is complete. This logic should bypass slots that are not populated or powered off.
20	RBT_ARB_IN	Output	NC-SI hardware arbitration input. If the baseboard supports multiple DPUs cards connected to the same RBT interface, it should implement logic that connects the RBT_ARB_IN pin of the first populated DPU card to its RBT_ARB_OUT pin if it is the only card present or to the RBT_ARB_OUT pin of the next populated card and so on sequentially for all cards on the specified RBT bus to ensure the arbitration ring is complete. This logic should bypass slots that are not populated or powered off.

4.4 Cabline CA-II Plus Connectors Pinouts

4.4.1 Component Side

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
1	GND	GND BAR		1
2	PCIE_REFCLK1_P	Micro coax	38	2

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
3	PCIE_REFCLK1_N	Micro coax	38	3
4	GND	GND BAR		4
5	PCIE_CPU_CX_15N	Micro coax	38	5
6	PCIE_CPU_CX_15P	Micro coax	38	6
7	GND	GND BAR		7
8	PCIE_CPU_CX_14N	Micro coax	38	8
9	PCIE_CPU_CX_14P	Micro coax	38	9
10	GND	GND BAR		10
11	PCIE_CPU_CX_13N	Micro coax	38	11
12	PCIE_CPU_CX_13P	Micro coax	38	12
13	GND	GND BAR		13
14	PCIE_CPU_CX_12N	Micro coax	38	14
15	PCIE_CPU_CX_12P	Micro coax	38	15
16	GND	GND BAR		16
17	PCIE_CPU_CX_11N	Micro coax	38	17
18	PCIE_CPU_CX_11P	Micro coax	38	18
19	GND	GND BAR		19
20	PCIE_CPU_CX_10N	Micro coax	38	20
21	PCIE_CPU_CX_10P	Micro coax	38	21
22	GND	GND BAR		22
23	PCIE_CPU_CX_9N	Micro coax	38	23
24	PCIE_CPU_CX_9P	Micro coax	38	24
25	GND	GND BAR		25
26	PCIE_CPU_CX_8N	Micro coax	38	26
27	PCIE_CPU_CX_8P	Micro coax	38	27
28	GND	GND BAR		28
29	PCIE_CPU_CX_7N	Micro coax	38	29
30	PCIE_CPU_CX_7P	Micro coax	38	30
31	GND	GND BAR		31
32	PCIE_CPU_CX_6N	Micro coax	38	32
33	PCIE_CPU_CX_6P	Micro coax	38	33
34	GND	GND BAR		34
35	PCIE_CPU_CX_5N	Micro coax	38	35
36	PCIE_CPU_CX_5P	Micro coax	38	36
37	GND	GND BAR		37
38	PCIE_CPU_CX_4N	Micro coax	38	38

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
39	PCIE_CPU_CX_4P	Micro coax	38	39
40	GND	GND BAR		40
41	PCIE_CPU_CX_3N	Micro coax	38	41
42	PCIE_CPU_CX_3P	Micro coax	38	42
43	GND	GND BAR		43
44	PCIE_CPU_CX_2N	Micro coax	38	44
45	PCIE_CPU_CX_2P	Micro coax	38	45
46	GND	GND BAR		46
47	PCIE_CPU_CX_1N	Micro coax	38	47
48	PCIE_CPU_CX_1P	Micro coax	38	48
49	GND	GND BAR		49
50	PCIE_CPU_CX_0N	Micro coax	38	50
51	PCIE_CPU_CX_0P	Micro coax	38	51
52	GND	GND BAR		52
53	I2C_DPU_BMC_SDA	Micro coax	38	53
54	I2C_DPU_BMC_SCL	Micro coax	38	54
55	AUX_PGOOD	Micro coax	38	55
56	No wire	Micro coax	38	56
57	I2C_AUX_SCL	Micro coax	38	57
58	I2C_AUX_SDA	Micro coax	38	58
59	S_PRST1_L	Micro coax	38	59
60	No wire			60

4.4.2 Print Side

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
1	SER_CLK	Micro coax	38	1
2	SER_CAPTURE	Micro coax	38	2
3	SER_DO	Micro coax	38	3
4	S_PERST2_CONN_L	Micro coax	38	4
5	SER_DI	Micro coax	38	5
6	Reserved_06	Micro coax	38	6
7	Reserved_07	Micro coax	38	7
8	Reserved_08	Micro coax	38	8
9	GND	GND BAR		9

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
10	PCIE_CX_CPU_0P	Micro coax	38	10
11	PCIE_CX_CPU_0N	Micro coax	38	11
12	GND	GND BAR		12
13	PCIE_CX_CPU_1P	Micro coax	38	13
14	PCIE_CX_CPU_1N	Micro coax	38	14
15	GND	GND BAR		15
16	PCIE_CX_CPU_2P	Micro coax	38	16
17	PCIE_CX_CPU_2N	Micro coax	38	17
18	GND	GND BAR		18
19	PCIE_CX_CPU_3P	Micro coax	38	19
20	PCIE_CX_CPU_3N	Micro coax	38	20
21	GND	GND BAR		21
22	PCIE_CX_CPU_4P	Micro coax	38	22
23	PCIE_CX_CPU_4N	Micro coax	38	23
24	GND	GND BAR		24
25	PCIE_CX_CPU_5P	Micro coax	38	25
26	PCIE_CX_CPU_5N	Micro coax	38	26
27	GND	GND BAR		27
28	PCIE_CX_CPU_6P	Micro coax	38	28
29	PCIE_CX_CPU_6N	Micro coax	38	29
30	GND	GND BAR		30
31	PCIE_CX_CPU_7P	Micro coax	38	31
32	PCIE_CX_CPU_7N	Micro coax	38	32
33	GND	GND BAR		33
34	PCIE_CX_CPU_8P	Micro coax	38	34
35	PCIE_CX_CPU_8N	Micro coax	38	35
36	GND	GND BAR		36
37	PCIE_CX_CPU_9P	Micro coax	38	37
38	PCIE_CX_CPU_9N	Micro coax	38	38
39	GND	GND BAR		39
40	PCIE_CX_CPU_10P	Micro coax	38	40
41	PCIE_CX_CPU_10N	Micro coax	38	41
42	GND	GND BAR		42
43	PCIE_CX_CPU_11P	Micro coax	38	43
44	PCIE_CX_CPU_11N	Micro coax	38	44
45	GND	GND BAR		45

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
46	PCIE_CX_CPU_12P	Micro coax	38	46
47	PCIE_CX_CPU_12N	Micro coax	38	47
48	GND	GND BAR		48
49	PCIE_CX_CPU_13P	Micro coax	38	49
50	PCIE_CX_CPU_13N	Micro coax	38	50
51	GND	GND BAR		51
52	PCIE_CX_CPU_14P	Micro coax	38	52
53	PCIE_CX_CPU_14N	Micro coax	38	53
54	GND	GND BAR		54
55	PCIE_CX_CPU_15P	Micro coax	38	55
56	PCIE_CX_CPU_15N	Micro coax	38	56
57	GND	GND BAR		57
58	S_PERST1_CONN_L	Micro coax	38	58
59	No wire	No Wire		59
60	S_PRST2_L	Micro coax	38	60

5 Hardware Installation and PCIe Bifurcation

Installation and initialization of the DPU/SuperNIC require attention to the mechanical attributes, power specification, and precautions for electronic equipment.

5.1 Safety Warnings

Safety warnings are provided here in the English language.

Please observe all safety warnings to avoid injury and prevent damage to system components. Note that not all warnings are relevant to all models.

Note that not all warnings are relevant to all models.

	General Installation Instructions Read all installation instructions before connecting the equipment to the power source.
	Jewelry Removal Warning Before you install or remove equipment that is connected to power lines, remove jewelry such as bracelets, necklaces, rings, watches, and so on. Metal objects heat up when connected to power and ground and can meltdown, causing serious burns and/or welding the metal object to the terminals.
	Over-temperature This equipment should not be operated in an area with an ambient temperature exceeding the maximum recommended: 55°C (131°F). An airflow of 200LFM at this maximum ambient temperature is required for HCA cards and NICs. To guarantee proper airflow, allow at least 8cm (3 inches) of clearance around the ventilation openings.
	During Lightning - Electrical Hazard During periods of lightning activity, do not work on the equipment or connect or disconnect cables.
	Copper Cable Connecting/Disconnecting Some copper cables are heavy and not flexible, as such, they should be carefully attached to or detached from the connectors. Refer to the cable manufacturer for special warnings and instructions.
	Equipment Installation This equipment should be installed, replaced, or serviced only by trained and qualified personnel.
	Equipment Disposal The disposal of this equipment should be in accordance to all national laws and regulations.
	Local and National Electrical Codes This equipment should be installed in compliance with local and national electrical codes.
	Hazardous Radiation Exposure • Caution - Use of controls or adjustment or performance of procedures other than those specified herein may result in hazardous radiation exposure. For products with optical ports. • CLASS 1 LASER PRODUCT and reference to the most recent laser standards: IEC 60 825-1:1993 + A1:1997 + A2:2001 and EN 60825-1:1994+A1:1996+ A2:20

5.2 Installation Procedure Overview

The installation procedure of DPU involves the following steps:

Step	Procedure	Direct Link
1	Check the system's requirements.	Refer to System Requirements
2	Pay attention to the airflow consideration within the host system	Refer to Airflow Requirements
3	Follow the safety precautions	Refer to Safety
4	Unpack the package	Refer to Unpacking
5	Follow the pre-installation checklist	Refer to Pre-Installation Checklist
7	Install the DPU according to the form-factor you have purchased.	Refer to DPU Installation
8	Connect cables or modules to the DPU	Refer to Cables and Modules
9	Power-up the DPU	Refer to DPU Power-Up Instructions

5.3 System Requirements

5.3.1 Hardware Requirements

Unless otherwise specified, products are designed to work in an environmentally controlled data center with low levels of gaseous and dust (particulate) contamination.

The operating environment should meet severity level G1 as per ISA 71.04 for gaseous contamination and ISO 14644-1 class 8 for cleanliness level.

The below table lists the motherboard and power supply requirements per DPU series.

OPNs	Power Supply Requirement
• B3140H SuperNICs: 900-9D3D4-00EN-HA0 and 900-9D3D4-00NN-HA0 • B3140L SuperNICs: 900-9D3B4-00EN-EA0 and 900-9D3B4-00PN-EA0 • B3220L SuperNICs: 900-9D3B4-00CV-EA0 and 900-9D3B4-00SV-EA0 • B3210L SuperNICs: 900-9D3B4-00CC-EA0 and 900-9D3B4-00SC-EA0	A minimum of 75W system power supply through the PCIe x16 interface
• B3220 DPUs: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0 • B3240 DPUs: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0 • B3210 DPUs: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0 • B3210E DPUs: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0	Require a supplementary 8-pin ATX power supply connectivity available through the external power supply connector.

5.3.2 Airflow Requirements

The DPU/SuperNIC is offered with one airflow direction: from the heatsink to the network ports.

Any use of the product in the opposite airflow direction (from network ports to heatsink) must be validated thermally to ensure proper cooling of the product.

Please refer to the [Specifications](#) section for airflow numbers per DPU model.

5.3.3 Software Requirements

- Refer to [System Requirements](#) section under the Introduction section.
- Software Stacks -BlueField-3 DPU/SuperNIC is shipped with Ubuntu - a Linux commercial operating system - which includes the NVIDIA OFED stack (MLNX_OFED), and is capable of running all customer-based Linux applications seamlessly. For more information, please refer to the DOCA SDK documentation or NVIDIA BlueField DPU BSP.

5.4 Safety Precautions

The DPU/SuperNIC is being installed in a system that operates with voltages that can be lethal. Before opening the case of the system, observe the following precautions to avoid injury and prevent damage to system components.

- Remove any metallic objects from your hands and wrists.
- Make sure to use only insulated tools.
- Verify that the system is powered off and is unplugged.
- It is strongly recommended to use an ESD strap or other antistatic devices.

5.5 Unpacking

Check against the package contents list that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping. Please note that the DPU/SuperNICs must be placed on an antistatic surface.

Please note that if the DPU is removed hastily from the antistatic bag, the plastic ziplock may harm the EMI fingers on the networking connector. Carefully remove the DPU/SuperNIC from the antistatic bag to avoid damaging the EMI fingers.

For package contents, please refer to [Package Contents](#).

5.6 Pre-Installation Checklist

1. Verify that your system meets the hardware and software requirements stated above.
2. Shut down your system if active.
Turn off the power to the system, and disconnect the power cord. Refer to the system

documentation for instructions. Before you install the DPU, make sure that the system is disconnected from power.

5.7 Installation Instructions

This section provides detailed instructions on how to install your DPU/SuperNIC in a system.

Choose the installation instructions according to the DPU configuration you would like to use.

OPNs	Installation Instructions
All DPUs	PCIe x16 Installation Instructions
• B3220 Model: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0 • B3240 Model: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0 • B3210 Model: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0 • B3210E Model: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0	[Optional] PCIe Extension Connection (2x PCIe x16) Installation Instructions

5.8 Cables and Modules

5.8.1 Networking Cable Installation

1. All cables can be inserted or removed with the unit powered on.
2. To insert a cable, press the connector into the port receptacle until the connector is firmly seated.
 - a. Support the weight of the cable before connecting the cable to the DPU/SuperNIC. Do this by using a cable holder or tying the cable to the rack.
 - b. Determine the correct orientation of the connector to the DPU/SuperNIC before inserting the connector. Do not try and insert the connector upside down. This may damage the DPU/SuperNIC.
 - c. Insert the connector into the DPU/SuperNIC . Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the DPU/SuperNIC.
 - d. Make sure that the connector locks in place.

When installing cables make sure that the latches engage.

Always install and remove cables by pushing or pulling the cable and connector in a straight line with the DPU.

3. After inserting a cable into a port, the Green LED indicator will light when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). See Networking Ports LEDs interface under the [Supported Interfaces](#) section.

4. After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When data is being transferred the Green LED will blink.
5. Make sure not to impede the air exhaust flow through the ventilation holes. Use cable lengths that allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.
6. To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. LED indicator will turn off when the cable is unseated.

5.8.2 8-pin ATX Power Supply Cable

The 8-pin ATX power supply cable is mandatory when powering-up the following DPUs. Without a connection to the power supply cable, the DPU will not complete the power on procedure and will not function properly.


- B3220 Model: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
- B3240 Model: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0
- B3210 Model: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- B3210E Model: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

The BlueField-3 DPU includes an 8-pin PCIe ATX power connector that provides additional power supply. While the PCIe slot feeds 66W, the 8-pin PCIe ATX power supply cable provides additional power.

Important Notes and Warnings

- The BlueField-3 DPU requires a standard PCIe ATX power connection. However, certain servers may require a custom setup to enable ATX power compatibility.
- Consult the manufacturer's manuals and specifications to ensure proper connectivity.
- Before connecting the ATX power cable, make sure you are using a PCIe-compliant 8-pin ATX cable from the server power supply kit.
- Do not link the CPU power cable to the BlueField-3 DPU PCIe ATX power connector, as their pin configurations differ. Using the CPU power cable in this manner is strictly prohibited and can potentially damage the BlueField-3 DPU. Please refer to [External PCIe Power Supply Connector Pins](#) for the external PCIe power supply pins.
- Avoid using non-standard cables that do not comply with the DPU, unnecessary adapter cables, or storing the cables near heat sources.
- It is preferable that the x16 PCIe golden fingers and the PCI ATX power supply draw from the same power source. For more information on how to power up the card, refer to [DPU Power-Up Instructions](#).
- The PCIe ATX 8-pin connector is not compatible with an EPS12V power cable source. Ensure that the appropriate PCIe auxiliary power source is available, not an EPS12V power source.
- If you are uncertain about your server's compatibility with the PCI ATX connection, please contact your NVIDIA representative for assistance.


Wiring Instructions



5.8.3 Cable Installation Instructions

- Ensure both the system and card are completely powered off for at least 20 seconds before inserting the power connector.
- DO NOT force the connector into place; the connectors are “keyed” to fit only one way. Refer to the mechanical keys below for further information.
- Connect the power supply end of the 8-pin connector to the appropriate receptor on the power supply unit.
- Apply parallel force when plugging the cable into the power supply connector. Do not bend or twist the cables when plugging or unplugging it.
- Ensure the cable is fully and securely connected to your DPU and the connector lock is secured.

Mechanical Keys of the 8-pin ATX Power Cable (PCIe CEM Specification Rev 5.0)




5.9 DPU Power-Up and Power-Down Sequences

The power-up and power-down sequences listed below apply to DPUs with x16 PCIe extension option.

- B3220 Model: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
- B3240 Model: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0
- B3210 Model: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- B3210E Model: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0


5.9.1 Power-Up Sequence

1. The 12V_ATX voltage can exceed the 12V_PCIE voltage by a maximum allowable voltage difference (dV) when the 12V_ATX reaches 11V. See below graph and table describing the dV between the 12V_ATX and 12V_PCIE voltages.




2. The 12V_ATX can be powered up after the 12V_PCIE, with a maximum delay of 2 seconds. The below graph illustrates the delay between the 12V_ATX and 12V_PCIE voltages at power-up.

Voltage



5.9.2 Power-Down Sequence

1. The 12V_PCIE voltage can be powered down simultaneously with the 12V_ATX voltage, or within a maximum delay of 200usec. The below graph illustrates the delay between the 12V_ATX and 12V_PCIE voltages at power-down.



- 2.
3. The 12V_PCIE voltage must not be powered down while the 12V_ATX voltage is powered up.


5.10 PCIe x16 Installation Instructions

5.10.1 Installation Instructions

This section provides detailed instructions on how to install your BlueField-3 card in a system.


Please note that the following figures are for illustration purposes only.

1. Open the system case.
2. Locate an available PCI Express slot.




3. Hold the card by its edges and bracket, avoiding contact with the top surface. Apply even pressure to both edges of the card and gently insert the BlueField-3 card into the PCI Express slot.


Exercise caution to avoid touching the components on the top surface of the card. Do not use excessive force when seating the card, as this may damage the system or the card's components.



4. When the BlueField-3 card is properly seated, the port connectors are aligned with the slot opening, and the card's faceplate is visible against the system chassis.



5. Secure the card with the screw.



6. Close the system case.
7. Install the networking cables. For instructions, please refer to [Networking Cable Installation](#).

5.10.2 Uninstalling the BlueField-3 Card

Safety Precautions

The card is installed in a system that operates with voltages that can be lethal. Before uninstalling the board, please observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. It is strongly recommended to use an ESD strap or other antistatic devices.
3. Turn off the system and disconnect the power cord from the server.

Card Removal

Please note that the following images are for illustration purposes only.

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention mechanism on the bracket (screws).
4. Holding the board from its center, gently pull the board out of the PCI Express slot.
5. When the port connectors reach the top of the chassis window, gently pull the board in parallel to the motherboard.

5.11 [Optional] PCIe Extension Connection (2x PCIe x16) Installation Instructions

The socket-direct functionality is currently not supported by firmware.

This section applies to the following DPUs when used as Socket Direct cards in dual-socket servers:

- B3220 DPUs: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
- B3240 DPUs: 900-9D3B6-00CN-AB0 900-9D3B6-00SN-AB0 and 900-9D3B6-00CN-PA0
- B3210 DPUs: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
- B3210E DPUs: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

For more information on the PCIe Auxiliary Kit, refer to [PCIe Auxiliary Card Kit](#).

The below images are for illustration purposes only.

The hardware installation section uses the terminology of white and black harnesses to differentiate between the two supplied cables. Due to supply chain variations, some DPUs may be supplied with two black harnesses instead. To clarify the difference between these two harnesses, one black harness was marked with a "WHITE" label and the other with a "BLACK" label.

The Cabline harness marked with a "WHITE" label should be connected to the connector on the DPU and Auxiliary PCIe card engraved with "White Cable", while the one marked with a "BLACK" label should be connected to the connector on the DPU and Auxiliary PCIe card engraved with "Black Cable".

The harnesses' minimal bending radius is 10[mm].

5.11.1 Installing the DPU

The installation instructions include steps that involve a retention clip to be used while connecting the Cabline harnesses to the DPUs. Please note that this is an optional accessory.


Please make sure to install the DPU cards in a PCIe slot that is capable of supplying the required power and airflow as stated in [Specifications](#).

➤ Connect the DPU with the Auxiliary connection card using the supplied Cabline CA-II Plus harnesses.


Step 1: Slide the black and white Cabline CA-II Plus harnesses through the retention clip while making sure the clip opening is facing the plugs.



Step 2: Plug the Cabline CA-II Plus harnesses on the DPU while paying attention to the color-coding. As indicated on both sides of the card; plug the black harness to the component side and the white harness to the print side.




Step 3: Verify the plugs are locked.



Step 4: Slide the retention clip latches through the cutouts on the PCB. The latches should face the annotation on the PCB.



Step 4: Clamp the retention clip. Verify both latches are firmly locked.




Verify that both latches are firmly snapped




Step 5: Slide the Cabline CA-II Plus harnesses through the retention clip. Make sure that the clip opening is facing the plugs.



Step 6: Plug the Cabline CA-II Plus harnesses on the PCIe Auxiliary Card. As indicated on both sides of the Auxiliary connection card; plug the black harness to the component side and the white harness to the print side.




Step 7: Verify the plugs are locked.




Step 8: Slide the retention clip through the cutouts on the PCB. Make sure latches are facing "Black Cable" annotation as seen in the below picture.



Step 9: Clamp the retention clip. Verify both latches are firmly locked.



Verify that both latches are firmly snapped



➤ Connect the DPU and PCIe Auxiliary Connection cards in available PCI Express x16 slots in the chassis.

Step 1: Locate two available PCI Express x16 slots.

Step 2: Applying even pressure at both corners of the cards, insert the DPU in the PCI Express slots until firmly seated.

Do not use excessive force when seating the cards, as this may damage the system or the cards.

Step 3: Applying even pressure at both corners of the cards, insert the Auxiliary Connection card in the PCI Express slots until firmly seated.

➤ Secure the DPU and PCIe Auxiliary Connection Cards to the chassis.

Secure the brackets to the chassis with the bracket screws.

➤ Install the PCIe external power cable.

Connect the 8-pin power connector from the power supply to the power connector on the top edge of the DPU. Note that the connector and socket on the graphics card have a unique shape and connect one way only. For further instructions, please refer to the cable vendor documentation. Please refer to the pinout description in [External Power Supply Connector](#).

➤ Close the system case.

5.11.2 Uninstalling the Cards

Safety Precautions

The DPU is installed in a system that operates with voltages that can be lethal. Before uninstalling the DPU, please observe the following precautions to avoid injury and prevent damage to system components.

1. Remove any metallic objects from your hands and wrists.
2. It is strongly recommended to use an ESD strap or other antistatic devices.
3. Turn off the system and disconnect the power cord from the server.

Card Removal

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention mechanisms on the brackets (clips or screws).
4. Holding the DPU from its center, gently pull the DPU and Auxiliary Connections card out of the PCI Express slot.

5.12 PCIe Bifurcation Configuration Options

PCIe bifurcation is supported starting from DOCA 2.5 with BlueField BSP 4.5.0 (released December 2023).


This section applies to the following OPNs:

B3220 DPUs: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0
B3240 DPUs: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0
B3210 DPUs: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0
B3210E DPUs: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0

NVIDIA BlueField-3 DPUs provide a range of configuration scenarios to meet the demands of environments and deployments. This section describes the various connectivity options for peripherals on the PCIe, including scenarios where the BlueField-3 DPU acts as the PCIe switch with NVMe SSDs as PCIe endpoints. While this list of scenarios is not exhaustive, it highlights the tested and verified options. Customers seeking to support unlisted configurations should contact NVIDIA Support.

The BlueField-3 DPU exposes two x16 PCIe interfaces, with internal PCIe switch architecture. The first interface is exposed via the x16 PCIe Gen 5.0/4.0 Goldfinger connector and serves as an endpoint to the host server by default. The additional PCIe x16 interface is exposed through the Cabline CA-II Plus connector, featuring programmable bifurcation as a downstream port. The following figure demonstrates the BlueField-3 DPU block diagram with the PCIe interfaces.

BlueField-3 DPU Block Diagram with PCIe Interfaces



The various configuration scenarios listed in this section include a diagram and instructions on how to bifurcate the PCIe interface using the mlxconfig tool. For more information on the mlxconfig tool, please refer to [mlxconfig - Changing Device Configuration Tool](#).

Before setting the desired configuration, take note of the following warnings:

- Any customer-set configuration overrides the previous configuration values.
- Misconfiguration may harm the system.
- It is recommended to establish out-of-band connectivity to the BlueField DPU Arm OS before setting any of these configurations for the first time. This enables you to reset the NVConfig parameters to their default values in case of misconfiguration.

The following table summarizes the available configuration scenarios.


Configuration	Root Port for Down Stream Port (DSP) devices: DPU ARM / Host	DPU PCIe Goldfingers Bifurcation	DPU PCIe Auxiliary Connection Bifurcation
Default	N/A	1 x Gen 5.0/4.0 x16 PCIe lanes as upstream port	1 x Gen 5.0/4.0 x16 PCIe lanes as upstream ports
Host as Root Port on Peripherals	Host	1 x Gen 5.0/4.0 x16 PCIe lanes as upstream port	4 x Gen 5.0/4.0 x4 PCIe lanes as downstream ports
DPU Arms as Root Port on Peripherals	DPU ARMs	1 x Gen 5.0/4.0 x16 PCIe lanes as upstream port	4 x Gen 5.0/4.0 x4 PCIe lanes as downstream ports

5.12.1 Host as Root Port on x4 PCIe Lane Peripherals

In this scenario, the x16 PCIe Goldfingers of the BlueField-3 DPU serve as endpoints to the host server (default), while the additional x16 PCIe lanes are accessible via the Cabline CA-II Plus connector, bifurcated into four PCIe links, where each link comprises x4 PCIe lanes.

In this configuration the host server assumes the role of the Root Port for downstream devices connected to the Cabline CA-II Plus connector. These downstream devices are exposed to the host server on its PCIe via the internal BlueField-3 DPU PCIe switch.

As seen in the below visual representation of this configuration, the host functions as the Root Port, branching into four PCIe links on the Cabline CA-II Plus connector, with each link featuring a bifurcation of x4 PCIe lanes.



Important Notes:

- mlxconfig can be configured either through the host in NIC Mode and DPU Mode, or directly from the DPU's Arm running OS.
- This configuration is persistent even following resets and NIC firmware updates.

The required set of configurations to implement this bifurcation is outlined below.

```
mlxconfig -d <device> s PCI_SWITCH0_UPSTRAEM_PORT_BUS=0
mlxconfig -d <device> s PCI_SWITCH0_UPSTRAEM_PORT_PEX=0
mlxconfig -d <device> s PCI_BUS00_HIERARCHY_TYPE=0
mlxconfig -d <device> s PCI_BUS00_WIDTH=5
mlxconfig -d <device> s PCI_BUS00_SPEED=4
mlxconfig -d <device> s PCI_BUS10_HIERARCHY_TYPE=1
mlxconfig -d <device> s PCI_BUS10_WIDTH=3
mlxconfig -d <device> s PCI_BUS10_SPEED=4
mlxconfig -d <device> s PCI_BUS12_HIERARCHY_TYPE=1
mlxconfig -d <device> s PCI_BUS12_WIDTH=3
```

```

mlxconfig -d <device> s PCI_BUS12_SPEED=4
mlxconfig -d <device> s PCI_BUS14_HIERARCHY_TYPE=1
mlxconfig -d <device> s PCI_BUS14_WIDTH=3
mlxconfig -d <device> s PCI_BUS14_SPEED=4
mlxconfig -d <device> s PCI_BUS16_HIERARCHY_TYPE=1
mlxconfig -d <device> s PCI_BUS16_WIDTH=3
mlxconfig -d <device> s PCI_BUS16_SPEED=4


```

5.12.2 DPU ARMs as Root Port on Peripherals

In this scenario, the x16 PCIe Gold Fingers of the BlueField-3 DPU serve as an endpoint to the host server (default), while the additional x16 PCIe lanes are accessible via the Cabline CA-II Plus connector, bifurcated into four PCIe links, where each link comprises x4 PCIe lanes.

In this configuration the DPU's Arm cores function as Root Port of the downstream devices connected to the Cabline CA-II Plus connector, and these remain unexposed to the host server on its PCIe.

As seen in the below visual representation of this configuration, the DPU ARMs operate as the Root Port, with bifurcation into four PCIe links on the Cabline CA-II Plus connector, where each link incorporates x4 PCIe lanes.



Important Notes:

- mlxconfig can be configured either through the host in NIC Mode and DPU Mode, or directly from the DPU's Arm running OS.
- This configuration is persistent even following resets and NIC firmware updates.

The required set of configurations to implement this bifurcation is outlined below:

```

mlxconfig -d <device> s PCI_BUS00_HIERARCHY_TYPE=0
mlxconfig -d <device> s PCI_BUS00_WIDTH=5
mlxconfig -d <device> s PCI_BUS00_SPEED=4

```

```
mlxconfig -d <device> s PCI_BUS10_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS10_WIDTH=3
mlxconfig -d <device> s PCI_BUS10_SPEED=4
mlxconfig -d <device> s PCI_BUS12_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS12_WIDTH=3
mlxconfig -d <device> s PCI_BUS12_SPEED=4
mlxconfig -d <device> s PCI_BUS14_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS14_WIDTH=3
mlxconfig -d <device> s PCI_BUS14_SPEED=4
mlxconfig -d <device> s PCI_BUS16_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS16_WIDTH=3
mlxconfig -d <device> s PCI_BUS16_SPEED=4
```

6 Setting High-Speed-Port Link Type

The following table lists the BlueField-3 supported speeds and the default networking port link type per OPN.

Model	OPN	Data Transmission Rate	Default Protocol and Rate
B3140L B3240	900-9D3B4-00EN-EA0 / 900-9D3B4-00PN-EA0 900-9D3B6-00CN-AB0 / 900-9D3B6-00SN-AB0 900-9D3B6-00CN-PA0	InfiniBand: NDR 400Gb/s Ethernet: 400GbE	InfiniBand NDR 400Gb/s
B3140H	900-9D3D4-00EN-HA0 / 900-9D3D4-00NN-HA0	InfiniBand: NDR 400Gb/s Ethernet: 400GbE	Ethernet 400GbE
B3220 B3220L	900-9D3B6-00CV-AA0 / 900-9D3B6-00SV-AA0 900-9D3B4-00CV-EA0 / 900-9D3B4-00SV-EA0	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE	Ethernet 200GbE
B3210 B3210E B3210L	900-9D3B6-00CC-AA0 / 900-9D3B6-00SC-AA0 900-9D3B6-00CC-EA0 / 900-9D3B6-00SC-EA0 900-9D3B4-00CC-EA0 / 900-9D3B4-00SC-EA0	InfiniBand: HDR100 100Gb/s Ethernet: 100GbE	Ethernet 100GbE

To configure the high-speed networking port mode, you can either use the [mlxconfig](#) or the [UEFI](#) tools.

UEFI can configure the device before the operating system is up, while mlxconfig configures the card once the operating system is up. According to your preference, use one of the below tools:

6.1 mlxconfig

The mlxconfig tool allows users to change device configurations without burning the firmware. The configuration is also kept after reset. By default, mlxconfig shows the configurations that will be loaded in the next boot. For more information and instructions, refer to [Using mlxconfig to Set IB/ETH Parameters](#).

6.2 UEFI

PreBoot drivers initialize the adapter device, check the port protocol type - Ethernet or InfiniBand - and bring up the port. Then it connects to a DHCP server to obtain its assigned IP address and network parameters and obtain the source location of the kernel/OS to boot from. The DHCP server instructs the PreBoot drivers to access the kernel/OS through a TFTP server, an iSCSI target, or some other service. For more information and instructions, refer to [UEFI](#).

7 Troubleshooting

Server unable to find the board	• Ensure that the board is placed correctly • Make sure the board slot and the board are compatible Install the board in a different PCI Express slot • Use the drivers that came with the board or download the latest • Make sure your motherboard has the latest BIOS • Try to reboot the server
The board no longer works	• Reseat the board in its slot or a different slot, if necessary • Try using another cable • Reinstall the drivers for the network driver files may be damaged or deleted • Reboot the server
boards stopped working after installing another board	• Try removing and re-installing all boards • Check that cables are connected properly • Make sure your motherboard has the latest BIOS
Link indicator light is off	• Try another port on the switch • Make sure the cable is securely attached • Check you are using the proper cables that do not exceed the recommended lengths • Verify that your switch and board port are compatible
Link light is on, but with no communication established	• Check that the latest driver is loaded • Check that both the board and its link are set to the same speed and duplex settings
Forgot password needed to install/upgrade the DPU image	Refer to the latest version of NVIDIA BlueField DPU BSP and follow instructions under "Updating DPU Software Packages" section.

8 Specifications

Ensure your system supports the following system hardware and power supply requirements prior to installing your card.

	Power Supply Requirement
• B3140H SuperNICs: 900-9D3D4-00EN-HA0 and 900-9D3D4-00NN-HA0 • B3140L SuperNICs: 900-9D3B4-00EN-EA0 and 900-9D3B4-00PN-EA0 • B3220L SuperNICs: 900-9D3B4-00CV-EA0 and 900-9D3B4-00SV-EA0 • B3210L SuperNICs: 900-9D3B4-00CC-EA0 and 900-9D3B4-00SC-EA0	A minimum of 75W system power supply through the PCIe x16 interface
• B3220 DPU: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0 • B3240 DPU: 900-9D3B6-00CN-AB0, 900-9D3B6-00SN-AB0 and 900-9D3B6-00CN-PA0 • B3210 DPU: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0 • B3210E DPU: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0	Require a supplementary 8-pin ATX power supply connectivity available through the external power supply connector.

8.1 B3140H SuperNICs Specifications

OPNs	• 900-9D3D4-00EN-HA0: Crypto Enabled with integrated BMC • 900-9D3D4-00NN-HA0: Crypto Disabled with integrated BMC	
BlueField-3 SoC	BlueField-3 E-Series -8 Arm-Cores - 505MHz/2000MHz	
Physical	Single-Slot HHHL Card Dimensions: 68.90mm x 167.65mm Tall Bracket Dimensions: 120.5 x 21.95mm Short Bracket Dimensions: 80.22mm x 22.83mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm	
Interfaces	See Supported Interfaces	
	On-board Memory	• Single-channel with 10 DDR5 + ECC (64bit + 8bit ECC) for a total of 16GB @ 5200MT/s • 40GB pSLC eMMC memory + 128GB SSD
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)
	Networking Connector	Single-port QSFP112 (copper and optical)
Data Rate	InfiniBand	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR
	Ethernet (Default)	400/200/100/50/25/10 Gb/s
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: NDR (4 lanes x 100Gb/s per lane) port, NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)	

OPNs	900-9D3D4-00EN-HA0: Crypto Enabled with integrated BMC 900-9D3D4-00NN-HA0: Crypto Disabled with integrated BMC		
	Ethernet: 400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage	12V	
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0° C to 55° C
		Non-operational	-40° C to 70° C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3D4	
Notes: (a) The BlueField-3 SuperNICs supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product. (b) The non-operational storage temperature specifications apply to the product without its package.			

8.2 B3140L SuperNICs Specifications

OPNs	900-9D3B4-00EN-EA0: Crypto Enabled with integrated BMC 900-9D3B4-00PN-EA0: Crypto Disabled with integrated BMC		
BlueField-3 SoC	BlueField-3 E-Series - 8 Arm-Cores - 505MHz/2000MHz		
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
Interfaces	See Supported Interfaces		
	On-board Memory	- Single-channel with 10 DDR5 + ECC (64bit + 8bit ECC) for a total of 16GB@ 5200MT/s 40GB pSLC eMMC memory + 128GB SSD	
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
	Networking Connector	Single-port QSFP112 (copper and optical)	

Data Rate	InfiniBand (Default Speed)	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet	400/200/100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: NDR (4 lanes x 100Gb/s per lane) port, NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Ethernet: 400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-2 C2M, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage	12V	
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B4	
Notes:			
(a) The BlueField-3 SuperNICs supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
(b) The non-operational storage temperature specifications apply to the product without its package.			

8.3 B3220L SuperNICs Specifications

OPNs	900-9D3B4-00CV-EA0: Crypto Enabled with integrated BMC 900-9D3B4-00SV-EA0: Crypto Disabled with integrated BMC
BlueField-3 SoC	BlueField-3 E-Series -8 Arm-Cores - 505MHz/2000MHz
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 120.5mm x 21.95mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm
Interfaces	See Supported Interfaces

OPNs	900-9D3B4-00CV-EA0: Crypto Enabled with integrated BMC 900-9D3B4-00SV-EA0: Crypto Disabled with integrated BMC		
	On-board Memory	- Single-channel with 10 DDR5 + ECC (64bit + 8bit ECC) for a total of 16GB @ 5200MT/s 40GB pSLC eMMC memory + 128GB SSD	
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
	Networking Connector	Dual-port QSFP112 (copper and optical)	
Data Rate	InfiniBand	NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet (Default)	200/100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Ethernet: 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage	12V	
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B4	
Notes:			
(a) The BlueField-3 SuperNICs supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
(b) The non-operational storage temperature specifications apply to the product without its package.			

8.4 B3210L SuperNICs Specifications

OPNs	900-9D3B4-00CC-EA0: Crypto Enabled with integrated BMC 900-9D3B4-00SC-EA0: Crypto Disabled with integrated BMC		
BlueField-3 SoC	BlueField-3 E-Series -8 Arm-Cores - 505MHz/2000MHz		
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 120.5mm x 21.95mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
Interfaces	See Supported Interfaces		
	On-board Memory	- Single-channel with 10 DDR5 + ECC (64bit + 8bit ECC) for a total of 16GB @ 5200MT/s 40GB pSLC eMMC memory + 128GB SSD	
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
	Networking Connector	Dual-port QSFP112 (copper and optical)	
Data Rate	InfiniBand	HDR100/EDR/FDR/SDR	
	Ethernet (Default)	100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Ethernet: 100GAUI-2 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2, 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M, XLPPI, SFI		
Electrical and Thermal Specifications	Voltage	12V	
	To be updated in a future revision of this document		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B4	
Notes:			
^(a) The BlueField-3 SuperNICs supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
^(b) The non-operational storage temperature specifications apply to the product without its package.			

8.5 B3240 DPUs Specifications

OPNs	900-9D3B6-00CN-AB0: Crypto Enabled with integrated BMC 900-9D3B6-00SN-AB0: Crypto Disabled with integrated BMC 900-9D3B6-00CN-PA0: Crypto Enabled with integrated BMC dedicated for cold-aisle servers only		
BlueField-3 IC	BlueField-3 P-Series - 16 Arm-Cores - 560MHz/2133MHz		
Physical	Dual-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 29.3mm		
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible) Optional: Additional PCIe x16 Gen 4.0/5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	On-board Memory	- Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5600MT/s 40GB pSLC eMMC memory + 128GB SSD	
	Networking Connector	Dual QSFP112 (copper and optical)	
Data Rate	InfiniBand (Default Speed)	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet	400/200/100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: NDR (4 lanes x 100Gb/s per lane) port, NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)		
	Ethernet: 400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-2 C2M, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
	PCI Express 5.0: SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)		
Electrical and Thermal Specifications	Voltage: 12V		
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	

Regulatory	Safety	CB / cTUVus / CE
	EMC	CE / FCC / VCCI / ICES / RCM
	RoHS	RoHS compliant
	Regulatory Model Number	D3B6
Notes: (a) The BlueField-3 DPU supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product. (b) The non-operational storage temperature specifications apply to the product without its package.		

8.6 B3240 DPU for Cold-Aisle Servers Specifications

OPNs	900-9D3B6-00CN-PA0: Crypto Enabled with integrated BMC dedicated for cold-aisle servers only	
BlueField-3 IC	BlueField-3 P-Series - 16 Arm-Cores - 560MHz/2133MHz	
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 29.3mm	
Interfaces	See Supported Interfaces	
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible) Optional: Additional PCIe x16 Gen 4.0/5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses
	On-board Memory	- Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5600MT/s 40GB pSLC eMMC memory + 128GB SSD
	Networking Connector	Dual QSFP112 (copper and optical)
Data Rate	InfiniBand (Default Speed)	NDR/NDR200/HDR/HDR100/EDR/FDR/SDR
	Ethernet	400/200/100/50/25/10 Gb/s
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: NDR (4 lanes x 100Gb/s per lane) port, NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane)	
	Ethernet: 400GAUI-4 C2M, 400GBASE-CR4, 200GAUI-2 C2M, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2, 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M, XLPPI, SFI	
	PCI Express 5.0: SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
Electrical and Thermal Specifications	Voltage: 12V	

	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B6	
Notes:			
(a) The BlueField-3 DPU supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
(b) The non-operational storage temperature specifications apply to the product without its package.			

8.7 B3210 DPUs Specifications

OPNs	900-9D3B6-00CC-AA0: Crypto Enabled with integrated BMC 900-9D3B6-00SC-AA0: Crypto Disabled with integrated BMC		
BlueField-3 SoC	BlueField-3 P-Series - 16 Arm-Cores - 560MHz/2133MHz		
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
		Optional: Additional PCIe x16 Gen 4.0/5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	On-board Memory	- Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5600MT/s 40GB pSLC eMMC memory + 128GB SSD	
Data Rate	Networking Connector	Dual-port QSFP112 (copper and optical)	
	InfiniBand	HDR100/EDR/FDR/SDR	
	Ethernet (Default Speed)	100/50/25/10 Gb/s	

Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane).		
	Ethernet: 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage: 12V		
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0° C to 55° C
		Non-operational	-40° C to 70° C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B6	
Notes:			
(a) The BlueField-3 DPU supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
(b) The non-operational storage temperature specifications apply to the product without its package.			

8.8 B3210E DPUs Specifications

Requires a supplementary 8-pin ATX power supply connectivity available through the external power supply connector.

OPNs	900-9D3B6-00CC-EA0: Crypto Enabled with integrated BMC 900-9D3B6-00SC-EA0: Crypto Disabled with integrated BMC		
BlueField-3 SoC	BlueField-3 E-Series - 16 Arm-Cores - 505MHz/2000MHz		
Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	

		Optional: Additional PCIe x16 Gen 4.0/5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	On-board Memory	- Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5200MT/s 40GB pSLC eMMC memory + 128GB SSD	
	Networking Connector	Dual-port QSFP112 (copper and optical)	
Data Rate	InfiniBand	HDR100/EDR/FDR/SDR	
	Ethernet (Default Speed)	100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(a) Auto-Negotiation: HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane).		
	Ethernet: 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage: 12V		
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0° C to 55° C
		Non-operational	-40° C to 70° C ^(b)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B6	
Notes:			
^(a) The BlueField-3 DPU supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.			
^(b) The non-operational storage temperature specifications apply to the product without its package.			

8.9 B3220 DPUs Specifications

OPNs	900-9D3B6-00CV-AA0: Crypto Enabled with integrated BMC 900-9D3B6-00SV-AA0: Crypto Disabled with integrated BMC
BlueField-3 SoC	BlueField-3 P-Series - 16 Arm-Cores - 560MHz/2133MHz

Physical	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
Interfaces	See Supported Interfaces		
	PCI Express Interface	Gen 5.0 SERDES @ 32.0GT/s, 16 lanes (4.0, 3.0, 2.0 and 1.1 compatible)	
		Optional: Additional PCIe x16 Gen 4.0/5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	On-Board Memory	- Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 32GB @ 5200MT/s^(a) 40GB pSLC eMMC memory + 128GB SSD	
	Networking Connector	Dual-port QSFP112 (copper and optical)	
Data Rate	InfiniBand	NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet (Default Speed)	200/100/50/25/10 Gb/s	
Protocol Support	InfiniBand: IBTA v1.5 ^(b) Auto-Negotiation: NDR200 (2 lanes x 100Gb/s per lane) port, HDR (50Gb/s per lane) port, HDR100 (2 lane x 50Gb/s per lane), EDR (25Gb/s per lane) port, FDR (14.0625Gb/s per lane), 1X/2X/4X SDR (2.5Gb/s per lane).		
	Ethernet: 200GAUI-2 C2M, 200GAUI-4 C2M, 200GBASE-CR4, 100GAUI-2 C2M, 100GAUI-1 C2M, 100GBASE-CR4, 100GBASE-CR2, 100GBASE-CR1, 50GAUI-2 C2M, 50GAUI-1 C2M, 50GBASE-CR, 50GBASE-R2 , 40GBASE-CR4, 40GBASE-R2, 25GBASE-R, 10GBASE-R, 10GBASE-CX4, 1000BASE-CX, CAUI-4 C2M, 25GAUI C2M, XLAUI C2M , XLPPI, SFI		
Electrical and Thermal Specifications	Voltage: 12V		
	Electrical and thermal specifications are provided in " NVIDIA BlueField-3 Networking Platform Product Specification " document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
Environmental	Temperature	Operational	0°C to 55°C
		Non-operational	-40°C to 70°C ^(c)
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety	CB / cTUVus / CE	
	EMC	CE / FCC / VCCI / ICES / RCM	
	RoHS	RoHS compliant	
	Regulatory Model Number	D3B6	




Notes:




- (c) 5600MT/s is supported from HW rev (AC). For more information, refer to the [latest NVIDIA BlueField DPU BSP](#).
- (b) The BlueField-3 DPU supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another NVIDIA InfiniBand product.
- (c) The non-operational storage temperature specifications apply to the product without its package.

8.10 DPUs Mechanical Drawing and Dimensions

All dimensions are in millimeters. The PCB mechanical tolerance is +/- 0.13mm.




The diagrams may differ for different cards and are provided here for illustration purposes only.




OPNs	Mechanical Drawings
HHHL Single-Slot SuperNICs Model: B3140H 900-9D3D4-00EN-HA0 900-9D3D4-00NN-HA0	
FHHL Single-Slot Single-Port SuperNICs Model: B3140L 900-9D3B4-00EN-EA0 900-9D3B4-00PN-EA0	
FHHL Single-Slot Dual-Port SuperNICs Model: B3220L 900-9D3B4-00CV-EA0 900-9D3B4-00SV-EA0 Model: B3210L 900-9D3B4-00CC-EA0 900-9D3B4-00SC-EA0	

FHHL Single-Slot Dual-Port DPUs with PCIe Extension Option Model: B3220 900-9D3B6-00CV-AA0 900-9D3B6-00SV-AA0 Model: B3210 900-9D3B6-00CC-AA0 900-9D3B6-00SC-AA0 Model: B3210E 900-9D3B6-00CC-EA0 900-9D3B6-00SC-EA0	
FHHL Dual-Slot Dual-Port DPUs Model: B3240 900-9D3B6-00CN-AB0 900-9D3B6-00SN-AB0	
FHHL Dual-Slot Dual-Port DPUs for Cold-Aisle Servers Only Model: B3240 900-9D3B6-00CN-PA0	

8.11 Bracket Mechanical Drawings

All dimensions are in millimeters.

DPU Configuration	Model/OPNs	Tall Bracket	Short Bracket
HHHL Single-Slot Dual-Port SuperNICs	Model: B3140H 900-9D3D4-00EN-HA0 900-9D3D4-00NN-HA0		
FHHL Single-Slot Single-port SuperNICs	Model: B3140L 900-9D3B4-00EN-EA0 900-9D3B4-00PN-EA0		

DPU Configuration	Model/OPNs	Tall Bracket	Short Bracket
FHHL Single-Slot Dual-Port SuperNICs	Model: B3210L 900-9D3B4-00CC-EA0 900-9D3B4-00SC-EA0 Model: B3220L 900-9D3B4-00CV-EA0 900-9D3B4-00SV-EA0		
FHHL Single-Slot Dual-Port DPUs	Model: B3220 900-9D3B6-00CV-AA0 900-9D3B6-00SV-AA0 Model: B3210 900-9D3B6-00CC-AA0 900-9D3B6-00SC-AA0 Model: B3210E 900-9D3B6-00CC-EA0 900-9D3B6-00SC-EA0 Model: B3240 for Cold-Aisle 900-9D3B6-00CN-PA0		
FHHL Dual-Slot Dual-port DPUs	Model: B3240 900-9D3B6-00CN-AB0 900-9D3B6-00SN-AB0 900-9D3B6-00CN-PA0		

9 Monitoring

9.1 Thermal Sensors

The thermal threshold definitions are described in the "*NVIDIA BlueField-3 Networking Platform Product Specification*" document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.

9.2 Heatsink

The heatsink is attached to the DPU by three screws to dissipate the heat from the SoC. The DPU SoC has a thermal shutdown safety mechanism that automatically shuts down the DPU in cases of high-temperature events, improper thermal coupling, or heatsink removal.

Refer to the table below for heatsink details per card configuration. For the required airflow (LFM) per OPN, please refer to [Specifications](#).

DPU/SmartNIC Configuration	OPN	Maximum Dimensions
Single-slot	900-9D3B4-00EN-EA0, 900-9D3B4-00PN-EA0, 900-9D3B6-00CV-AA0, 900-9D3B6-00SV-AA0, 900-9D3B6-00CC-AA0, 900-9D3B6-00SC-AA0, 900-9D3B6-00CC-EA0, 900-9D3B6-00SC-EA0, 900-9D3B6-00CN-PA0	Length, Width, Height: 139.6mm x 92.7mm x 10.2mm
Dual-slot	900-9D3B6-00CN-AB0, 900-9D3B6-00SN-AB0	Length, Width, Height: 139.6mm x 92.7mm x 29.3mm

10 Finding the GUID/MAC on the Board

Each BlueField-3 platform is uniquely identified to facilitate efficient management and configuration. This section aims to guide you through the process of identifying your BlueField-3 platform by using the information provided on the board label.


The board label includes unique identifiers essential for network communication— specifically, a primary GUID address for InfiniBand protocols or/and MAC addresses for Ethernet protocols. If your BlueField-3 board supports both InfiniBand and Ethernet protocols, the label will feature a primary GUID along with five ‘base’ MAC addresses derived from the GUID address: Host, ECPF, MPF, DPU BMC, and OOB.

The table below lists the different BlueField-3 form factors. Refer to the relevant section depending on the board you have purchased. For each form factor, we provide an example of the board label and a table listing the different MAC addresses, offering examples and additional information to enhance clarity.

OPNs	Direct Link
• B3220 DPUs: 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0 • B3240 DPUs: 900-9D3B6-00CN-AB0 and 900-9D3B6-00SN-AB0 • B3210 DPUs: 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0 • B3210E DPUs: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0	DPUs Board Label Example
• B3140H SuperNICs: 900-9D3D4-00EN-HA0 and 900-9D3D4-00NN-HA0 • B3140L SuperNICs: 900-9D3B4-00EN-EA0 and 900-9D3B4-00PN-EA0 • B3220L SuperNICs: 900-9D3B4-00CV-EA0 and 900-9D3B4-00SV-EA0 • B3210L SuperNICs: 900-9D3B4-00CC-EA0 and 900-9D3B4-00SC-EA0	SuperNICs Board Label Example

10.1 DPUs Board Label Example

Note: The product revisions (“Rev”) indicated on the label do not necessarily represent the latest revisions of the cards.




Item on Label	Description	Example/Additional Information
GUID (Global Unique Identifier)	A unique 16 byte identifier of the BlueField in an InfiniBand network. The GUID identifier never changes.	The GUID on the label is assigned to the System GUID of the BlueField HCA. The System GUID represents the HCA hardware with all its ports. Each HCA port is assigned a dedicated Node GUID. The Node GUID of port 0 is equal to the System GUID. The Node GUID of port 1 is equal to the System GUID+1
HOST (Base MAC)	Represents the base MAC address for high-speed data port 0.	DPU's with more than one high-speed port will use incremental (+1) MAC addresses against the base HOST MAC on the label. For instance: From the label - HOST: 94 6D AE F5 A1 CC Port 0: 94 6D AE F5 A1 CC (same as label) Port 1: 94 6D AE F5 A1 CD (+1 from base) Port 2: 94 6D AE F5 A1 CE (+2 from base) Port 3: 94 6D AE F5 A1 CF (+3 from base)
ECPF (Embedded CPU Physical Function)	Includes the embedded Arm system responsible for managing NIC resources and controlling the data path.	Each physical high-speed port is allocated one ECPF. The base ECPF MAC is displayed on the label and is associated with Port 0 . The second port (Port 1) receives base-ECPF+1 MAC address and so forth, similar to HOST MAC allocation pattern described above.
MPF (Multi/Management Physical Function)	For future use.	N/A
OOB (Out-of-Band)	Indicates the MAC address of the 1GbE out-of-band management port within the DPU's embedded Arm system.	Network traffic for the "OOB" (DPU embedded Arm system) and "DPU BMC" traverse through the 1GbE RJ-45 connector located on the front panel of the DPU card.

Item on Label	Description	Example/Additional Information
DPU BMC (Baseboard Management Controller)	Represents the MAC address assigned to the 1GbE out-of-band management port of the DPU BMC. Note: The base MAC address for the DPU BMC is one higher (+1) than the OOB MAC address.	
S/N (Serial Number)	A unique serial number that distinguishes your DPU from other boards. This S/N is used for tracking and support purposes.	N/A

10.2 SuperNICs Board Label Example

Note: The product revisions (“Rev”) indicated on the label do not necessarily represent the latest revisions of the cards.



Item on Label	Description	Example/Additional Information
GUID (Global Unique Identifier)	A unique 16 byte identifier of the BlueField in an InfiniBand network. The GUID identifier never changes.	The GUID on the label is assigned to the System GUID of the BlueField SuperNIC. The System GUID represents the HCA hardware with all its ports. Each HCA port is assigned a dedicated Node GUID. The Node GUID of port 0 is equal to the System GUID. The Node GUID of port 1 is equal to the System GUID+1

Item on Label	Description	Example/Additional Information
HOST (Base MAC)	Represents the base MAC address for high-speed data port 0.	SuperNICs with more than one high-speed port use incremental (+1) MAC addresses against the base HOST MAC on the label. For instance: From the label - HOST: 94 6D AE F5 A1 CC Port 0: 94 6D AE F5 A1 CC (same as label) Port 1: 94 6D AE F5 A1 CD (+1 from base) Port 2: 94 6D AE F5 A1 CE (+2 from base) Port 3: 94 6D AE F5 A1 CF (+3 from base)
ECPF (Embedded CPU Physical Function)	Includes the embedded Arm system responsible for managing NIC resources and controlling the data path.	Each physical high-speed port is allocated one ECPF. The base ECPF MAC is displayed on the label and is associated with Port 0 . The second port (Port 1) receives base-ECPF+1 MAC address and so forth, similar to HOST MAC allocation pattern described above.
MPF (Multi/Management Physical Function)	For future use.	N/A
OOB (Out-of-Band)	Indicates the MAC address of the 1GbE out-of-band management port within the SuperNIC's embedded Arm system.	Network traffic for the "OOB" (DPU embedded Arm system) and "DPU BMC" traverse through the 1GbE RJ-45 connector located on the front panel of the SuperNIC.
DPU BMC (Baseboard Management Controller)	Represents the MAC address assigned to the 1GbE out-of-band management port of the DPU BMC. Note: The base MAC address for the DPU BMC is one higher (+1) than the OOB MAC address.	
S/N (Serial Number)	A unique serial number that sets your DPU apart from other boards. This S/N is used for tracking and support purposes.	N/A

11 PCIe Auxiliary Card Kit

This section applies to the following DPUs:

- 900-9D3B6-00CN-AB0
- 900-9D3B6-00SN-AB0
- 900-9D3B6-00CV-AA0
- 900-9D3B6-00SV-AA0

There are two available PCIe auxiliary kit functionalities:

1. Utilizing the Socket-Direct capability, where the PCIe extension card is connected to the BlueField-3 DPU, used as an end-point.
2. Utilizing the Down Stream Port (DSP) extension option, where the PCIe extension card is connected to the BlueField-3 DPU, used as a root complex for storage devices.

The two Cabline SA-II Plus harnesses in the PCIe auxiliary kit have different routings. To distinguish between these two harnesses, one black harness is marked with a “WHITE” label while the harness is marked with a “BLACK” label.

The Cabline harness marked with the “WHITE” label should be connected to the connector on the networking card and PCIe Auxiliary card engraved with “White Cable” while the one marked with the “BLACK” label should be connected to the connector on the networking card and the PCIe Auxiliary card engraved with “Black Cable”. The Cabline SA-II Plus harness mates with two 60-pin connectors (P/N 20790-060E-01), on both sides. The black Cabline SA-II Plus harness mates with the connector on the component side (top side) of the network card, while the White Cabline SA-II Plus harness mates with the print side (bottom side) of the main network card. For hardware installation, please refer to [\[Optional\] PCIe Extension Connection \(2x PCIe x16\) Installation Instructions](#).

11.1 Socket Direct

The socket-direct functionality is currently not supported by firmware. Please approach your sales representatives.

Socket Direct network cards, which cost-effectively integrate a single network adapter silicon on a primary board, and an auxiliary PCIe connection card and Cabline SA-II Plus Harness connecting the two. Socket Direct enables direct access from each CPU to the network through its dedicated PCIe interface as the card's 32-lane PCIe bus is split into two 16-lane buses, with one bus accessible through a PCIe x16 edge connector and the other bus through an x16 Auxiliary PCIe Connection card. The two cards should be installed into two PCIe x16 slots and connected using two Cabline SA-II Plus harnesses.

The PCIe auxiliary kit can be purchased separately to operate in a dual-socket server. The table below lists the available PCIe auxiliary kit ordering part numbers, depending on the desired length of the Cabline SA-II Plus harnesses and the PCI Express interface, Gen 4.0 or Gen 5.0.

Ordering Part Number	Passive Auxiliary Connection	Cabline SA-II Plus Harnesses Length
MTMK9100-T15	PCIe Gen 5.0 x16 connection card	2x 150mm harnesses
MTMK9100-T25	PCIe Gen 4.0 x16 connection card	2x 250mm harnesses
MTMK9100-T35	PCIe Gen 4.0 x16 connection card	2x 350mm harnesses

11.2 Down Stream Port (DSP)

The BlueField-3 DPUs with downstream port extension option integrate a single network adapter silicon on a primary board, and an auxiliary PCIe connection card and [Cabline SA-II Plus Harness](#) connecting the two.

This enables direct access from each CPU to the network through its dedicated PCIe interface as the card's 32-lane PCIe bus is split into two 16-lane buses, with one bus accessible through a PCIe x16 edge connector and the other bus through an x16 Auxiliary PCIe Connection card. The two cards should be installed into two PCIe x16 slots and connected using two Cabline SA-II Plus harnesses.

The PCIe auxiliary kit can be purchased separately to allow downstream port extension. The table below lists the available PCIe auxiliary kit ordering part numbers, depending on the desired length of the [Cabline SA-II Plus harnesses](#) and the PCI Express interface, Gen 4.0 or Gen 5.0.

OPN	Description
930-9DAX5-0015-000	NVIDIA BlueField-3 Auxiliary Kit for additional self-hosted PCIe Gen 5.0/4.0 x16 connection, two 150mm Cabline SA-II Plus Harnesses
930-9DAX4-0035-000	NVIDIA BlueField-3 Auxiliary Kit for additional self-hosted PCIe Gen 4.0 x16 connection, two 350mm Cabline SA-II Plus Harnesses
930-9DHAR-0001-000	NVIDIA BlueField-3 150mm Gen 5.0 FRU harness, Cabline SA-II Plus 60-pin micro coax 38AWG

PCIe Auxiliary Card Package Contents

Category	Qty	Item
Cards	1	For Socket Direct: MTKM9100-T15: PCIe x16 Gen 5.0 Auxiliary Connection Card MTKM9100-T25 and MTKM9100-T35: PCIe x16 Gen 4.0 Auxiliary Connection Card For Down Stream Port (DSP): 930-9DAX5-0015-000: PCIe x16 Gen 5.0/4.0 Auxiliary Connection Card 930-9DAX4-0035-000: PCIe x16 Gen 4.0 Auxiliary Connection Card
Harnesses	1	Cabline CA-II Plus harness (white) - Length according to kit OPN
	1	Cabline CA-II Plus harness (black) - Length according to kit OPN

Category	Qty	Item
Accessories	2	Retention Clip for Cabeline harness (shipped assembled on the harnesses - optional)
	1	PCIe Auxiliary card short bracket
	1	PCIe Auxiliary card tall bracket (shipped assembled on the Auxiliary card)

11.3 Channel Insertion Loss

Channel insertion loss is the signal power loss resulting from a device's insertion in a transmission line or optical fiber and is usually expressed in decibels (dB).

The following table describes the NVIDIA® BlueField®-3 channel insertion loss budget for PCIe Gen 5.0 architecture (32 GT/s).

The total PCIe channel insertion loss approved by PCI-SIG Gen5.0 spec is 36dB @16GHz.

The total BlueField-3 DPU board insertion loss of the PCIe lanes (PCORE1) routed to the Cabline CA-II Plus is 6dB (@16GHz).

The Passive Socket Direct PCIe Auxiliary Card Loss is 1.5dB (@16Ghz).

The Cabline CA-II Plus harness loss at 16GHz:

Harness Length	Channel Loss at Gen 5.0
15cm	3.8dB
35cm	7.6dB
55cm	11.4dB

The above is measured data; it is recommended to add 0.5dB margins for your system (some loss variations are possible).

The Cabline CA-II Plus harnesses loss = 0.24dB/cm for Gen 5.0.

The above loss includes the Cabline CA-II Plus harnesses and connectors on both sides.

The PCI-SIG Gen5 SPEC also defines the total loss for AIC (bump to GF) to be 9.5dB @16Ghz.

Please note that the BlueField-3 AIC, together with a 15cm Cabline CA-II Plus harnesses and the Passive PCIe Auxiliary Card loss is 10.3dB, which is approved by NVIDIA.

11.4 Cabline CA-II Plus Harness Pinouts

11.4.1 Cabline CA-II Plus Harness - Component Side

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
1	GND	GND BAR			1

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
2	PCIE_REFCLK1_P	Micro coax	Primary PCIe clock from the motherboard to the BlueField DPU Main card, to be used for the x16 Cabline harness PCIe interface. This clock must meet all the PCIe SIG spec requirements. It should be driven from the motherboard side.	38	2
3	PCIE_REFCLK1_N	Micro coax	Primary PCIe clock from the motherboard to the BlueField DPU Main card, to be used for the x16 Cabline harness PCIe interface. This clock must meet all the PCIe SIG spec requirements. It should be driven from the motherboard side.	38	3
4	GND	GND BAR			4
5	PCIE_CPU_CX_15N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	5
6	PCIE_CPU_CX_15P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	6
7	GND	GND BAR			7
8	PCIE_CPU_CX_14N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	8

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
9	PCIE_CPU_CX_14P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	9
10	GND	GND BAR			10
11	PCIE_CPU_CX_13N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	11
12	PCIE_CPU_CX_13P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	12
13	GND	GND BAR			13
14	PCIE_CPU_CX_12N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	14
15	PCIE_CPU_CX_12P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	15
16	GND	GND BAR			16
17	PCIE_CPU_CX_11N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	17
18	PCIE_CPU_CX_11P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	18

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
19	GND	GND BAR			19
20	PCIE_CPU_CX_10N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	20
21	PCIE_CPU_CX_10P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	21
22	GND	GND BAR			22
23	PCIE_CPU_CX_9N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	23
24	PCIE_CPU_CX_9P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	24
25	GND	GND BAR			25
26	PCIE_CPU_CX_8N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	26
27	PCIE_CPU_CX_8P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	27
28	GND	GND BAR			28

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
29	PCIE_CPU_CX_7N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	29
30	PCIE_CPU_CX_7P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	30
31	GND	GND BAR			31
32	PCIE_CPU_CX_6N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	32
33	PCIE_CPU_CX_6P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	33
34	GND	GND BAR			34
35	PCIE_CPU_CX_5N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	35
36	PCIE_CPU_CX_5P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	36
37	GND	GND BAR			37
38	PCIE_CPU_CX_4N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	38

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
39	PCIE_CPU_CX_4P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	39
40	GND	GND BAR			40
41	PCIE_CPU_CX_3N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	41
42	PCIE_CPU_CX_3P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	42
43	GND	GND BAR			43
44	PCIE_CPU_CX_2N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	44
45	PCIE_CPU_CX_2P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	45
46	GND	GND BAR			46
47	PCIE_CPU_CX_1N	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	47
48	PCIE_CPU_CX_1P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/ GPU/End-Point	38	48

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
49	GND	GND BAR			49
50	PCIE_CPU_CX_ON	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/GPU/End-Point	38	50
51	PCIE_CPU_CX_0P	Micro coax	Connect from the Black Cabline CA-II Plus cable through 220nF capacitors to the x16 PCIe Tx lanes of the CPU/GPU/End-Point	38	51
52	GND	GND BAR			52
53	I2C_DPU_BMC_SDA	Micro coax		38	53
54	I2C_DPU_BMC_SCL	Micro coax		38	54
55	AUX_PGOOD	Micro coax		38	55
56	No wire	Micro coax		38	56
57	I2C_AUX_SCL	Micro coax	The BlueField silicon serves as the I2C bus master on this bus. An I2C EEPROM at I2C address 0x57 needs to be mounted on the motherboard side to report to the Cabline CA-II Plus interface parameters to the main-card BlueField DPU silicon, like Cabline CA-II Plus cables length (contact NVIDIA for the format of this EEPROM). If additional optional I2C slave devices need to be managed by the main-card BlueField DPU silicon, they need to be included on this I2C bus as well.	38	57

Pin#	Signal Name	Wire Type	Detailed Description	AWG#	Pin# on the other end
58	I2C_AUX_SDA	Micro coax	The BlueField silicon serves as the I2C bus master on this bus. An I2C EEPROM at I2C address 0x57 needs to be mounted on the motherboard side to report to the Cabline CA-II Plus interface parameters to the main-card BlueField DPU silicon, like Cabline CA-II Plus cables length (contact NVIDIA for the format of this EEPROM). If additional optional I2C slave devices need to be managed by the main-card BlueField DPU silicon, they need to be included on this I2C bus as well.	38	58
59	S_PRSNT1_L	Micro coax	Connect this pin to GND No wires are connected to these pins to ensure they do not interfere with the operation of S_PRSNT2_L for the detection when the two Cabline harnesses are installed.	38	59
60		No Wire			60

11.4.2 Cabline CA-II Plus Harness - Print Side

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
1	SER_CLK	Micro coax	38	This pin is used as the serializer clock (SER_CLK) from the DPU to the device/s located on the motherboard.	1

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
2	SER_CAPTURE	Micro coax	38	This pin is used as the serializer capture (SER_CAPTURE).	2
3	SER_DO	Micro coax	38	This pin is used as the serializer data out from the BlueField DPU to the device/s located on the motherboard.	3
4	S_PERST2_CONN_L	Micro coax	38	Optional: PCIe compliant PERST_L (active low PCI Reset) signal for the Cabline CA-II Plus PCIe interface. To be used as the PERST_L signal for the control of PCIe lane 15:8, when a bifurcation of the Cabline CA-II Plus PCIe x16 interface to two x8 interfaces is needed (and in specific main board assemblies which support such bifurcation). The direction of this optional PERST_L signal depends on the implementation: When connecting a CPU root complex to the Cabline CA-II Plus PCIe interface, this signal is driven from the motherboard side (from the CPU), to the BlueField DPU; When connecting a GPU or an end point to the Cabline CA-II Plus PCIe interface, this signal is driven from the BlueField DPU (which operates as a PCIe switch in this case), to the GPU or end-point on the motherboard side.	4

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
5	SER_DI	Micro coax	38	This pin is used as the serializer data in from the device/s located on the motherboard.	5
6	Reserved_06	Micro coax	38		6
7	Reserved_07	Micro coax	38	Reserved for future expansion	7
8	Reserved_08	Micro coax	38	Reserved for future expansion	8
9	GND	GND BAR			9
10	PCIE_CX_CPU_0P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	10
11	PCIE_CX_CPU_0N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	11
12	GND	GND BAR			12
13	PCIE_CX_CPU_1P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	13
14	PCIE_CX_CPU_1N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	14
15	GND	GND BAR			15

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
16	PCIE_CX_CPU_2P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	16
17	PCIE_CX_CPU_2N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	17
18	GND	GND BAR			18
19	PCIE_CX_CPU_3P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	19
20	PCIE_CX_CPU_3N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	20
21	GND	GND BAR			21
22	PCIE_CX_CPU_4P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	22

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
23	PCIE_CX_CPU_4N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	23
24	GND	GND BAR			24
25	PCIE_CX_CPU_5P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	25
26	PCIE_CX_CPU_5N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	26
27	GND	GND BAR			27
28	PCIE_CX_CPU_6P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	28
29	PCIE_CX_CPU_6N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	29
30	GND	GND BAR			30

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
31	PCIE_X_CPU_7P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	31
32	PCIE_CX_CPU_7N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	32
33	GND	GND BAR			33
34	PCIE_CX_CPU_8P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	34
35	PCIE_CX_CPU_8N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	35
36	GND	GND BAR			36
37	PCIE_CX_CPU_9P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	37

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
38	PCIE_CX_CPU_9N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	38
39	GND	GND BAR			39
40	PCIE_CX_CPU_10P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	40
41	PCIE_CX_CPU_10N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	41
42	GND	GND BAR			42
43	PCIE_CX_CPU_11P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	43
44	PCIE_CX_CPU_11N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	44
45	GND	GND BAR			45

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
46	PCIE_CX_CPU_12P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	46
47	PCIE_CX_CPU_12N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	47
48	GND	GND BAR			48
49	PCIE_CX_CPU_13P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	49
50	PCIE_CX_CPU_13N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	50
51	GND	GND BAR			51
52	PCIE_CX_CPU_14P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	52

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
53	PCIE_CX_CPU_14N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	53
54	GND	GND BAR			54
55	PCIE_CX_CPU_15P	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	55
56	PCIE_CX_CPU_15N	Micro coax	38	Connect from the White Cabline CA-II Plus harness to the PCIe x16 Rx lanes of the CPU/GPU/End-Point. 220nF caps are required on this signal on the BlueField DPU.	56
57	GND	GND BAR			57

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
58	S_PERST1_CONN_L	Micro coax	38	PCIe compliant PERST_L (active low PCI Reset) signal for the PCIe Cabline CA-II Plus Connectors. The direction of this PERST_L signal depends on the implementation: When connecting a CPU root complex to the PCIe Cabline CA-II Plus interface, this signal is driven from the motherboard side (from the CPU), to the BlueField DPU. When connecting a GPU or an end point to the PCIe Cabline CA-II Plus interface, this signal is driven from the BlueField DPU side (which operates as a PCIe switch in this case), to the GPU or end-point on the motherboard side.	58
59		No Wire			59

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
60	S_PRSNT2_L	Micro coax	38	Connect to a 4.7K pull-up resistor to 3.3V on the motherboard side, to detect if both the Cabline harnesses are connected or not. This signal is connected to S_PRSNT1_L on the BlueField DPU. In the motherboard side, read logic low if both Cabline harnesses are connected. Read logic 1 (3.3V) if one or both the Cabline harnesses are not connected. No wires are connected to these pins to ensure they do not interfere with the operation of S_PRSNT1_L for the detection when the two Cabline harnesses are installed.	60


11.5 Technical Specifications

Physical	PCIe Auxiliary Card Size: 5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black) Length: 15, 25 or 35cm		
Power Consumption	Voltage: 12V, 3.3V_AUX Maximum current: 100mA for the 3.3V_AUX voltage rail		
PCIe Connectivity	MTMK9100-T15	PCI Express Gen 5.0/4.0: SERDES @ 16/32 GT/s, x16 lanes (Gen 3.0 compatible)	
	MTMK9100-T25 / MTMK9100-T35	PCI Express Gen 4.0: SERDES @ 16GT/s, x16 lanes (Gen 3.0 compatible)	
Environmental	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C
	Humidity	Operational	10% to 85% relative humidity
		Non-operational	10% to 90% relative humidity
	Altitude (Operational)	3050m	
Regulatory	Safety: CB / cTUVus / CE		

Physical	PCIe Auxiliary Card Size: 5.09 in. x 2.32 in. (129.30mm x 59.00mm) Two Cabline CA-II Plus harnesses (white and black) Length: 15, 25 or 35cm
Power Consumption	Voltage: 12V, 3.3V_AUX Maximum current: 100mA for the 3.3V_AUX voltage rail
	EMC: CE / FCC / VCCI / ICES / RCM / KC
	RoHS: RoHS Compliant

11.5.1 PCIe Auxiliary Card Mechanical Drawings and Dimensions


All dimensions are in millimeters. The PCB mechanical tolerance is +/- 0.13mm.



11.5.2 Bracket Mechanical Drawings and Dimensions

Auxiliary PCIe Connection Card Tall Bracket	Auxiliary PCIe Connection Card Short Bracket
A side-view mechanical drawing of the Auxiliary PCIe Connection Card Tall Bracket. It is a grey, elongated component with a series of rectangular slots. Green dimension lines indicate a length of 121.05mm and a height of 21.59mm.	A side-view mechanical drawing of the Auxiliary PCIe Connection Card Short Bracket. It is a grey, shorter component with a series of rectangular slots. Green dimension lines indicate a length of 80.29mm and a height of 22.83mm.

11.5.3 Cabline CA-II Plus Harnesses Mechanical Drawing



12 Supported Servers and Power Cords

12.1 Supported Servers

Server support depends on the particular setup being used. The following is a partial list of servers with which the DPUs have been tested. For more information, please contact your NVIDIA representative.

Dell	Lenovo	HPE	Supermicro	IEI (Inspur)	H3C	Gigabyte	ASUS	Aivres
R760xa XE9680	SR675v3	DL385 Gen11 XD670 DL380 Gen11	SYS-421GE-TNRT SYS-521GE-TNRT SYS-821GE-TNHR	NF5468M7-T NF5688A7 NF5688M7 NF5468A7	R5300 G6 R5500CG6 R5500G6	G493-SB0-A G593-SD2-A G593-SD1-A G593-ZD0-L G593-ZD1-A G593-ZD2-A G493-ZB0-A	ESC N8-E11 ESC N8A-E11 ESC8000-E11/ E11P ESC8000A-E12/E12P ESC4000-E11 ESC4000A-E12	KR6288X2 KR6288E2

12.2 Supported Power Cords

Vendor	Server	Cable Vendor	Part Number	Description
SuperMicro	SYS-421GE-TNRT	Supermicro	CBL-PWEX-1148-20 CBL-PWEX-1040	Supermicro 20cm 8-Pin to 8-Pin Power Cable Supermicro 8-pin to Two 6+2 Pin 5cm Power GPU Cable
	SYS-521GE-TNRT			
	SYS-821GE-TNHR	Supermicro	CBL-PWEX-0665	Supermicro 8-Pin CPU to 8-Pin PCIe 30cm GPU Power Cable
Lenovo	SR675v3	Molex	SC17B44959	2x4 Micro-Hi ST to Molex 2x4 (455860005) ST 18 AWG L=740 mm (SVT)

13 Document Revision History

Date	Description
Jul 2024	- Added wiring diagram to the 8-pin ATX Power Supply Cable installation instructions - Added warnings to the eMMC and NVMe SSD Interfaces in Supported Interfaces
Jun 2024	- Added 900-9D3B6-00CN-PA0 support across the document - Updated Supported Servers and Power Cords
Mar 2024	- Updated B3210E PCIe speed to Gen 5.0 - Removed the "Securing Workloads" feature from Features and Benefits - Updated the onboard memory specifications in Features and Benefits - Updated NC-SI trance lengths in NC-SI Management Interface - Updated the important notes and warnings in External PCIe Power Supply Connector - Added 8-pin ATX Power Supply Cable installation instructions - Updated Cabline CA-II Plus Connectors Pinouts - Updated PCIe Auxiliary Card Kit - Added notes across the document stating that the Socket-Direct functionality is currently not supported by firmware - Updated Supported Servers and Power Cords - Updated the DDR5 specifications for the B3220L and B3220 platforms across the document
Feb 2024	Updated Supported Servers
Jan 2023	- Fixed a typo in Specifications - Updated PCIe Gen support for B3210E DPUs
Dec 2023	- Updated the document to include BlueField-3 SuperNICs upon reaching the "Mass Production" milestone. OPNs: B3210L, B3220L and B3210E - Updated the eMMC interface in Features and Benefits - Added PCIe Bifurcation Configuration Options - Updated Finding the GUID/MAC on the Board
Nov 2023	- Updated Channel Insertion Loss - Added regulatory label in Finding the GUID/MAC on the Board - Updated Supported Servers
Oct 2023	Updated the following sections: DPU Power-Up Instructions. - Added NC-SI trance lengths in NC-SI Management Interface. - Updated PCIe Auxiliary Card Kit to indicate that MTMK9100-T15 supports Gen 5.0. - Updated NVMe SSD Interface. - Updated default speed of B3140H to 400Gb/s Ethernet. - Added Regulatory Model Number for each DPU in Specifications.
Aug 2023	- Added step 3 to section Verifying DPU Connection and Setting Up Host Environment - Fixed typo in Bracket Mechanical Drawings - Updated board label in Finding the GUID/MAC on the Board
Aug 2023	- Updated the lifecycle tag of 900-9D3B6-00CV-AA0 and 900-9D3B6-00SV-AA0 to indicate "Mass Production". - Updated Cabline CA-II Plus Connectors Pinouts. - Updated Cabline CA-II Plus Harness Pinouts. - Added a note to NVMe SSD Interface.
July 2023	Amended a note on airflow direction in Hardware Installation and PCIe Bifurcation.

Date	Description
June 2023	- Added <code>sudo</code> to step 2 in section Verifying DPU Connection and Setting Up Host Environment. - Updated list of identified devices in section Verifying DPU Connection and Setting Up Host Environment. - Added section Connecting to BlueField and Verifying Version. - Updated step 2.c. in section Updating BlueField BFB Image. - Added <code>sudo</code> to step 3.b. in section Updating BlueField BFB Image. - Marked 900-9D3B6-00CC-AA0 and 900-9D3B6-00SC-AA0 as EOL (End of Life) products. - Added new DPUs to the user manual: 900-9D3B6-00CC-EA0 and 900-9D3B6-00SC-EA0. - Updated Cabline CA-II Plus Harness Pinouts. - Updated SoC frequency for E-Series DPUs in Specifications
May 2023	- Updated Specifications - added non-operational storage temperature specifications. - Updated Ethernet protocols in Specifications.
Apr 2023	Added PSID and device ID information in NVIDIA BlueField-3 Networking Platform User Guide .
Mar 2023	Updated DDR5 SDRAM On-Board Memory .
Feb 2023	Updated External PCIe Power Supply Connector .
Feb 2022	- Updated list of SKUs across the document - Added BlueField-3 Administrator Quick Start Guide - Added Setting High-Speed-Port Link Type - Added an important note on the External PCIe Power Supply Connector
Jan 2023	Added PCIe Auxiliary Card Kit
Nov 2022	Updated the following sections: NC-SI Interface Pinouts
Jul 2022	Updated the following sections: Cabline CA-II Plus Connectors with additional information. - Added Cabline CA-II Connector pins in Pinouts Description. Finding the GUID/MAC on the Board with board label examples. PCI Express Interface pinouts to reflect changes in pins BB81 and B882. - Added heatsink dimensions in Introduction and Specifications.
Jun 2022	Renamed the document from "NVIDIA BlueField-3 InfiniBand/VPI DPU User Guide" to "NVIDIA BlueField-3 DPU User Guide"
May 2022	First release

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. Neither NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: “NVIDIA”) make any representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice. Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of NVIDIA Corporation and/or



Mellanox Technologies Ltd. in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2024 NVIDIA Corporation & affiliates. All Rights Reserved.

