



# **NVIDIA BlueField-3 DPU Controller User Manual**

# Table of Contents

<b>1</b>	<b>Introduction.....</b>	<b>7</b>
1.1	System Requirements .....	7
1.2	Package Contents .....	8
1.2.1	Card Package .....	8
1.2.2	Accessories Kit .....	8
1.2.3	PCIe Auxiliary Card Package.....	8
1.3	Features and Benefits .....	8
<b>2</b>	<b>BlueField-3 Administrator Quick Start Guide .....</b>	<b>13</b>
2.1	Prerequisites for Initial BlueField-3 Deployment .....	13
2.2	First-time Installation Procedure.....	13
<b>3</b>	<b>Supported Interfaces .....</b>	<b>14</b>
3.1	FHHL DPU Controller Layout and Interface Information .....	14
3.2	Interfaces Detailed Description.....	15
3.2.1	System-on-Chip (SoC) .....	15
3.2.2	Networking Interface.....	16
3.2.3	Networking Ports LEDs Interface .....	16
3.2.4	PCI Express Interface.....	17
3.2.5	DDR5 SDRAM On-Board Memory.....	17
3.2.6	NC-SI Management Interface .....	17
3.2.7	UART Interface Connectivity .....	17
3.2.8	USB 4-pin RA Connector.....	18
3.2.9	1GbE OOB Management Interface .....	18
3.2.10	PPS IN/OUT Interface .....	19
3.2.11	External PCIe Power Supply Connector.....	19
3.2.12	Cabline CA-II Plus Connectors.....	20
3.2.13	Integrated BMC Interface .....	21
3.2.14	NVMe SSD Interface.....	21
3.2.15	RTC Battery .....	21
3.2.16	eMMC Interface.....	21
<b>4</b>	<b>Pinouts Description .....</b>	<b>22</b>
4.1	PCI Express Interface.....	22
4.2	External Power Supply Connector .....	24

4.3	NC-SI Management Interface .....	25
4.4	Cabline CA-II Plus Connectors Pinouts .....	27
4.4.1	Component Side.....	27
4.4.2	Print Side .....	29
<b>5</b>	<b>Hardware Installation and PCIe Bifurcation .....</b>	<b>32</b>
5.1	Safety Warnings .....	32
5.2	Installation Procedure Overview.....	33
5.3	System Requirements .....	33
5.3.1	Hardware Requirements .....	33
5.3.2	Airflow Requirements .....	33
5.3.3	Software Requirements .....	34
5.4	Safety Precautions .....	34
5.5	Unpacking .....	34
5.6	Pre-Installation Checklist .....	34
5.7	Installation Instructions.....	34
5.8	Cables and Modules.....	35
5.8.1	Networking Cable Installation .....	35
5.8.2	8-pin ATX Power Supply Cable .....	35
5.8.3	Cable Installation Instructions .....	36
5.9	DPU Controller Power-Up and Power-Down Sequences .....	37
5.9.1	Power-Up Sequence .....	37
5.9.2	Power-Down Sequence .....	38
5.10	PCIe x16 DPU Controller Installation Instructions.....	39
5.10.1	Installation Instructions.....	39
5.10.2	Uninstalling the BlueField-3 Card .....	40
5.11	[Optional] PCIe Extension Connection (2x PCIe x16) Installation Instructions ....	41
5.11.1	Installing the DPU.....	42
5.11.2	Uninstalling the Cards.....	46
5.12	PCIe Bifurcation Configuration Options .....	47
5.12.1	Bifurcated to Eight x4 PCIe.....	48
5.12.2	Bifurcated to One x16 PCIe and Eight x2 PCIe .....	49
<b>6</b>	<b>Setting High-Speed-Port Link Type .....</b>	<b>51</b>
6.1	mlxconfig.....	51
6.2	UEFI .....	51

<b>7</b>	<b>Troubleshooting .....</b>	<b>52</b>
7.1	General Troubleshooting .....	52
<b>8</b>	<b>Specifications.....</b>	<b>53</b>
8.1	900-9D3C6-00CV-GA0 / 900-9D3C6-00CV-DA0 Specifications.....	53
8.1.1	DPU Controller Mechanical Drawing and Dimensions.....	54
<b>9</b>	<b>Monitoring.....</b>	<b>55</b>
9.1	Thermal Sensors.....	55
9.2	Heatsink .....	55
<b>10</b>	<b>Finding the GUID/MAC and Serial Number on the Card .....</b>	<b>56</b>
10.1	DPU Controller Board Label Example .....	56
<b>11</b>	<b>PCIe Auxiliary Card Kit .....</b>	<b>58</b>
11.1	PCIe Auxiliary Card Package Contents .....	58
11.2	Channel Insertion Loss .....	59
11.3	Cabline CA-II Plus Harness Pinouts .....	59
11.3.1	Cabline CA-II Plus Harness - Component Side .....	59
11.3.2	Cabline CA-II Plus Harness - Print Side .....	66
11.4	Technical Specifications .....	75
11.4.1	PCIe Auxiliary Card Mechanical Drawings and Dimensions.....	76
11.4.2	Bracket Mechanical Drawings and Dimensions .....	76
11.4.3	Cabline CA-II Plus Harnesses Mechanical Drawing .....	77
<b>12</b>	<b>Cables and Cabling Configurations .....</b>	<b>78</b>
12.1	RJ45 Cable for the 1GbE OOB Management Interface.....	78
12.2	USB Interface Cabling .....	78
<b>13</b>	<b>General Cooling Recommendations for Heatsink-less DPU Controllers .....</b>	<b>79</b>
13.1	General Notes .....	79
13.2	Components that Require Cooling.....	80
13.2.1	Mechanical Interfaces .....	81
13.2.2	Encapsulation and Battery on SFG .....	81
<b>14</b>	<b>Document Revision History .....</b>	<b>83</b>

## About This Manual

This User Manual describes NVIDIA® BlueField®-3 BF3220SH DPU Controller cards. It provides details as to the interfaces of the board, specifications, required software and firmware for operating the card, hardware installation, driver installation and bring-up instructions.

## Ordering Part Numbers

The below table provides the ordering part numbers (OPN) for the available BlueField BF3220SH DPU Controllers.

Model	NVIDIA SKU		Heatsink	Lifecycle
B3220SH	<b>900-9D3C6-00CV-GA0</b>	NVIDIA BlueField-3 B3220SH E-Series <b>No heatsink</b> FHHL Storage Controller, 200GbE (default mode) / NDR200 IB, Dual-port QSFP112, PCIe Gen5.0 x16 with x16 PCIe extension option, 16 Arm cores, 48GB on-board DDR, integrated BMC, Crypto Enabled, Secure Boot	No Heatsink	Prototype
	<b>900-9D3C6-00CV-DA0</b>	NVIDIA BlueField-3 B3220SH E-Series FHHL Storage Controller, 200GbE (default mode) / NDR200 IB, Dual-port QSFP112, PCIe Gen5.0 x16 with x16 PCIe extension option, 16 Arm cores, 48GB on-board DDR, integrated BMC, Crypto Enabled, Secure Boot	✓	Prototype

## 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](http://www.nvidia.com) → Support
- E-mail: [enterprisesupport@nvidia.com](mailto: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

### Related Documentation

<a href="#">InfiniBand Architecture Specification</a>	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.
<a href="#">IEEE Std 802.3 Specification</a>	IEEE Ethernet specification.

<a href="#"><u>InfiniBand Architecture Specification</u></a>	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.
<a href="#"><u>PCI Express Specifications</u></a>	Industry Standard PCI Express Base and Card Electromechanical Specifications.
<a href="#"><u>NVIDIA LinkX Interconnect Solutions</u></a>	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.
<a href="#"><u>BlueField DPU Platform BSP Documentation</u></a>	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.
<a href="#"><u>DOCA SDK Software Documentation</u></a>	NVIDIA DOCA SDK software.

#### Document Conventions

When discussing memory sizes, GB and GBytes are used in this document to mean size in giga Bytes. The use of Gb or Gbits (small b) indicates size in giga bits. 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 data processing unit (DPU) Controller is the 3rd-generation data center infrastructure-on-a-chip that enables organizations to build software-defined, hardware-accelerated IT infrastructures from cloud to core data center to edge. With 200Gb/s Ethernet or NDR200 200Gb/s InfiniBand network connectivity, BlueField-3 DPU Controller offloads, accelerates, and isolates software-defined networking, storage, security, and management functions in ways that profoundly improve data center performance, efficiency, and security. Providing powerful computing, and a broad range of programmable acceleration engines in the I/O path, BlueField-3 is perfectly positioned to address the infrastructure needs of the most demanding applications, while delivering full software backward compatibility through the NVIDIA DOCA™ software framework.

BlueField-3 DPU Controllers are highly optimized for NVMe storage systems, Network Functions Virtualization (NFV), Cloud and Machine Learning workloads. BlueField-3 integrates all the discrete components of a storage system appliance into a single chip making it the premier solution for building Just-a-Bunch-Of-Flash (JBOF) systems, All-Flash-Array and storage appliances for NVMe over Fabrics. With an integrated NVMe-oF and NVMe/TCP offload accelerator, the BlueField-3 DPU Controller has a superior performance advantage over existing JBOF systems, significantly reducing storage transaction latency, while increasing IOPs (I/O operations per second).

## 1.1 System Requirements

<b>Motherboard</b>	<b>In PCIe x16 Configuration</b> PCIe Gen 5.0 (32GT/s) through x16 edge connector. <b>In PCIe x16 Extension Option - Switch DSP (Data Stream Port)</b> <ul style="list-style-type: none"><li>• PCIe Gen 5.0 SERDES @32GT/s through edge connector</li><li>• PCIe Gen 5.0 SERDES @32GT/s through PCIe Auxiliary Connection Card</li></ul>
<b>System Power Supply</b>	The maximum power consumption is up to <b>150W</b> and is split between the two power sources as follows: <ul style="list-style-type: none"><li>• Up to 66W from the PCIe golden fingers (12V)</li><li>• The rest of the consumed power is drawn from a supplementary 8-pin ATX power supply connector</li></ul> <div style="border: 1px solid #f9e79f; padding: 5px; margin-top: 10px;">NOTE: The power cable is not included in the package. It should be part of the system wiring or ordered separately as an accessory.</div>
<b>Operating System</b>	The BlueField-3 DPU Controller 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-3 Software User Manual.
<b>Connectivity</b>	<ul style="list-style-type: none"><li>• Interoperable with 1/10/25/40/50/100/200 Gb/s Ethernet switches and SDR/DDR/EDR/HDR100/HDR/NDR200 InfiniBand switches</li><li>• Passive copper cable with ESD protection</li><li>• Powered connectors for optical and active cable support</li></ul>

## 1.2 Package Contents

Prior to unpacking your DPU Controller, 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

<b>Cards</b>	1x BlueField-3 DPU Controller
<b>Accessories</b>	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.

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

### 1.2.3 PCIe Auxiliary Card Package

This is an optional kit.

The PCIe auxiliary kit can be purchased separately to operate selected DPU Controllers in a dual-socket server. For package contents, refer to PCIe Auxiliary Card Kit.

For more detailed information, refer to Specifications.

## 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																															
<b>InfiniBand Architecture Specification v1.5 compliant</b>	<p>The DPU Controller delivers low latency, high bandwidth, and computing efficiency for high-performance computing (HPC), artificial intelligence (AI), and hyperscale cloud data centers applications.</p> <p>The DPU Controller is InfiniBand Architecture Specification v1.5 compliant.</p> <p><b>InfiniBand Network Protocols and Rates:</b></p> <table border="1" data-bbox="671 461 1390 1249"> <thead> <tr> <th data-bbox="671 461 810 607" rowspan="2">Protocol</th> <th data-bbox="810 461 967 607" rowspan="2">Standard</th> <th colspan="2" data-bbox="967 461 1283 519">Rate (Gb/s)</th> <th data-bbox="1283 461 1390 607" rowspan="2">Comments</th> </tr> <tr> <th data-bbox="967 519 1123 607">4x Port (4 Lanes)</th> <th data-bbox="1123 519 1283 607">2x Ports (2 Lanes)</th> </tr> </thead> <tbody> <tr> <td data-bbox="671 607 810 797">NDR/NDR200</td> <td data-bbox="810 607 967 797">IBTA Vol2 1.5</td> <td data-bbox="967 607 1123 797">425</td> <td data-bbox="1123 607 1283 797">212.5</td> <td data-bbox="1283 607 1390 797">PAM4 256b/257b encoding and RS-FEC</td> </tr> <tr> <td data-bbox="671 797 810 987">HDR/HDR100</td> <td data-bbox="810 797 967 987">IBTA Vol2 1.4</td> <td data-bbox="967 797 1123 987">212.5</td> <td data-bbox="1123 797 1283 987">106.25</td> <td data-bbox="1283 797 1390 987">PAM4 256b/257b encoding and RS-FEC</td> </tr> <tr> <td data-bbox="671 987 810 1122">EDR</td> <td data-bbox="810 987 967 1122">IBTA Vol2 1.3.1</td> <td data-bbox="967 987 1123 1122">103.125</td> <td data-bbox="1123 987 1283 1122">51.5625</td> <td data-bbox="1283 987 1390 1122">NRZ 64b/66b encoding</td> </tr> <tr> <td data-bbox="671 1122 810 1249">FDR</td> <td data-bbox="810 1122 967 1249">IBTA Vol2 1.2</td> <td data-bbox="967 1122 1123 1249">56.25</td> <td data-bbox="1123 1122 1283 1249">N/A</td> <td data-bbox="1283 1122 1390 1249">NRZ 64b/66b encoding</td> </tr> </tbody> </table>					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
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																								
<b>Up to 200 Gigabit Ethernet</b>	The DPU Controller complies with the following IEEE 802.3 standards: 200GbE / 100GbE / 50GbE / 40GbE / 25GbE / 10GbE																								
	<table border="1"> <thead> <tr> <th data-bbox="671 351 1026 405">Protocol</th> <th data-bbox="1026 351 1390 405">MAC Rate</th> </tr> </thead> <tbody> <tr> <td data-bbox="671 405 1026 483">IEEE802.3ck</td> <td data-bbox="1026 405 1390 483">200/100 Gigabit Ethernet (Include ETC enhancement)</td> </tr> <tr> <td data-bbox="671 483 1026 645">IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu</td> <td data-bbox="1026 483 1390 645">200/100 Gigabit Ethernet (Include ETC enhancement)</td> </tr> <tr> <td data-bbox="671 645 1026 723">IEEE 802.3bj IEEE 802.3bm</td> <td data-bbox="1026 645 1390 723">100 Gigabit Ethernet</td> </tr> <tr> <td data-bbox="671 723 1026 824">IEEE 802.3by Ethernet Technology Consortium</td> <td data-bbox="1026 723 1390 824">25/50 Gigabit Ethernet</td> </tr> <tr> <td data-bbox="671 824 1026 869">IEEE 802.3ba</td> <td data-bbox="1026 824 1390 869">40 Gigabit Ethernet</td> </tr> <tr> <td data-bbox="671 869 1026 913">IEEE 802.3ae</td> <td data-bbox="1026 869 1390 913">10 Gigabit Ethernet</td> </tr> <tr> <td data-bbox="671 913 1026 1025">IEEE 802.3cb</td> <td data-bbox="1026 913 1390 1025">2.5/5 Gigabit Ethernet (For 2.5: support only 2.5 x1000BASE-X)</td> </tr> <tr> <td data-bbox="671 1025 1026 1104">IEEE 802.3ap</td> <td data-bbox="1026 1025 1390 1104">Based on auto-negotiation and KR startup</td> </tr> <tr> <td data-bbox="671 1104 1026 1171">IEEE 802.3ad IEEE 802.1AX</td> <td data-bbox="1026 1104 1390 1171">Link Aggregation</td> </tr> <tr> <td data-bbox="671 1171 1026 1283">IEEE 802.1Q IEEE 802.1P VLAN tags and priority</td> <td data-bbox="1026 1171 1390 1283"></td> </tr> <tr> <td data-bbox="671 1283 1026 1525">IEEE 802.1Qau (QCN) Congestion Notification IEEE 802.1Qaz (ETS) IEEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)</td> <td data-bbox="1026 1283 1390 1525"></td> </tr> </tbody> </table>	Protocol	MAC Rate	IEEE802.3ck	200/100 Gigabit Ethernet (Include ETC enhancement)	IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu	200/100 Gigabit Ethernet (Include ETC enhancement)	IEEE 802.3bj IEEE 802.3bm	100 Gigabit Ethernet	IEEE 802.3by Ethernet Technology Consortium	25/50 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) IEEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)	
	Protocol	MAC Rate																							
	IEEE802.3ck	200/100 Gigabit Ethernet (Include ETC enhancement)																							
	IEEE802.3cd IEEE802.3bs IEEE802.3cm IEEE802.3cn IEEE802.3cu	200/100 Gigabit Ethernet (Include ETC enhancement)																							
	IEEE 802.3bj IEEE 802.3bm	100 Gigabit Ethernet																							
	IEEE 802.3by Ethernet Technology Consortium	25/50 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) IEEE 802.1Qbb (PFC) IEEE 802.1Qbg IEEE 1588v2 IEEE 802.1AE (MACSec) Jumbo frame support (9.6KB)																								

Feature	Description
<b>On-board Memory</b>	<ul style="list-style-type: none"> <li>• Quad SPI NOR FLASH - includes 256Mbit for Firmware image.</li> <li>• UVPS EEPROM - includes 2Mbit.</li> <li>• 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.</li> <li>• DPU_BMC Flashes: <ul style="list-style-type: none"> <li>• 2x 64MByte for BMC Image</li> <li>• 512MByte for Config Data</li> </ul> </li> <li>• eMMC pSLC 40GB with 30K Write Cycles eMMC for SoC BIOS and OS.</li> <li>• SSD (onboard BGA) 128GByte for user SoC OS, logs and application SW.</li> <li>• DDR5 SDRAM - 48GB @5200MT/s dual-channel DDR5 SDRAM memory. Solder down on-board 128bit + 16bit ECC.</li> </ul>
<b>BlueField-3 SoC</b>	<p>The DPU Controller integrates 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.</p>
<b>RDMA and RDMA over Converged InfiniBand/Ethernet (RoCE)</b>	<p>The DPU Controller, utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged InfiniBand/Ethernet) technology, delivers 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.</p>
<b>Storage Acceleration</b>	<ul style="list-style-type: none"> <li>• 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</li> <li>• The DPU Controller 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, the DPU Controller 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.</li> </ul>
<b>NVMe-oF and NVMe/TCP Accelerations</b>	<p>Non-volatile Memory Express (NVMe) over Fabrics or over TCP are protocols for communicating block storage IO requests over RDMA or TCP to transfer data between a host computer and a target solid-state storage device or system over a network. The DPU Controller may operate as a co-processor offloading specific storage tasks from the host using its powerful NVMe Offload accelerator.</p>
<b>SR-IOV</b>	<p>The DPU Controller SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.</p>

Feature	Description
<b>GPU Direct</b>	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. The DPU Controller uses high-speed DMA transfers to copy data between P2P devices resulting in more efficient system applications
<b>Cryptography Accelerations</b>	From IPsec and TLS data-in-motion inline encryption to AES-XTS block-level data-at-rest encryption and public key acceleration, the DPU Controller 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.
<b>Security Accelerators</b>	A consolidated compute and network solution based on the DPU Controller achieves significant advantages over a centralized security server solution. Standard encryption protocols and security applications can leverage NVIDIA BlueField-3 compute capabilities and network offloads for security application solutions such as Layer4 Statefull Firewall.
<b>Out-of-Band Management</b>	The DPU Controller 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 the status of other in-band network components.
<b>BMC</b>	The DPU Controller 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 DPU boards.

---

## 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

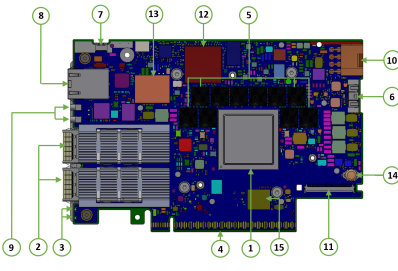
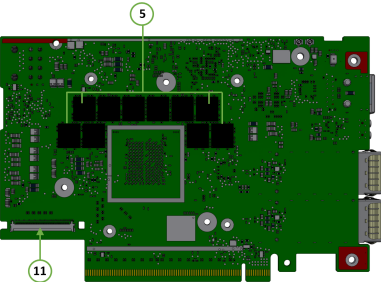
Unable to render include or excerpt-include. Could not retrieve page.

## 3 Supported Interfaces

This section describes the DPU Controller supported interfaces. Each numbered interface referenced in the figures is described in the following table with a link to detailed information.

The below figures are for illustration purposes only and might not reflect the current revision of the DPU Controller.

### 3.1 FHHL DPU Controller Layout and Interface Information

OPN	DPU Controller Component Side	DPU Controller Print Side
900-9D3C6-00CV-GA0 900-9D3C6-00CV-DA0		

Item	Interface	Description
1	<a href="#">System-on-Chip (SoC)</a>	16 Arm-Cores SoC
2	<a href="#">Networking Interface</a>	The network traffic is transmitted through the DPU Controller QSFP112 connectors. The QSFP112 connectors allow the use of modules and optical and passive cable interconnect solutions
3	<a href="#">Networking Ports LEDs Interface</a>	One bi-color I/O LEDs <b>per port</b> to indicate link and physical status
4	<a href="#">PCI Express Interface</a>	PCIe Gen 5.0 through an x16 edge connector
5	<a href="#">DDR5 SDRAM On-Board Memory</a>	20 units of DDR5 SDRAM for a total of 48GB @ 5200MT/s. 128bit + 16bit ECC, solder-down memory
6	<a href="#">NC-SI Management Interface</a>	NC-SI 20 pins BMC connectivity for remote management
7	<a href="#">USB 4-pin RA Connector</a>	Used for OS image loading
8	<a href="#">1GbE OOB Management Interface</a>	1GbE BASE-T OOB management interface

Item	Interface	Description
9	<a href="#">MMCX RA PPS IN/OUT</a>	Allows PPS IN/OUT
10	<a href="#">External PCIe Power Supply Connector</a>	An external 12V power connection through an 8-pin ATX connector
11	<a href="#">Cabline CA-II Plus Connectors</a>	Two Cabline CA-II plus connectors are populated to allow connectivity to an additional PCIe x16 Auxiliary card
12	<a href="#">Integrated BMC</a>	DPU BMC
13	<a href="#">SSD Interface</a>	128GB
14	<a href="#">RTC Battery</a>	Battery holder for RTC
15	<a href="#">eMMC</a>	x8 NAND flash

## 3.2 Interfaces Detailed Description

### 3.2.1 System-on-Chip (SoC)

NVIDIA BlueField-3 DPU Controller is a family of advanced DPU Controller 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 DPU Controller IC 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.

#### 3.2.1.1 Encryption


DPU Controller addresses the concerns of modern data centers by combining hardware encryption accelerators with embedded software and fully integrated advanced network capabilities, making it an ideal platform for developing proprietary security applications. It enables a distributed security architecture by isolating and protecting each workload and providing 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. DPU Controller can deliver powerful functionality, including encryption of data-in-motion, bare-metal provisioning, stateful L4 firewall, and more.

### 3.2.2 Networking Interface

The DPU Controller includes 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.

### 3.2.3 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		
Error	4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:		
	<b>Error Type</b>	<b>Description</b>	<b>LED Behavior</b>
	I <sup>2</sup> C	I <sup>2</sup> 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		



### 3.2.4 PCI Express Interface

The DPU Controller 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

### 3.2.5 DDR5 SDRAM On-Board Memory

The DPU Controllers incorporate 20 units of DDR5 SDRAM for a total of 48GB @ 5200MT/s single or dual channels. 64/128bit + 8bit ECC per channel, solder-down memory.

### 3.2.6 NC-SI Management Interface

The DPU Controller 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 RMI and has a connector on the DPU Controller. Please refer to [NC-SI Management Interface](#) for pins.

The maximum trace length on the board is 5.7inch.

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

### 3.2.7 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.

### 3.2.8 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 DPU Controller card box and should be ordered separately as part of the accessories kit (P/N: MBF35-DKIT).

### 3.2.9 1GbE OOB Management Interface

The DPU Controller 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 DPU Controllers with integrated BMC: 1GbE OOB Management can be performed via the integrated BMC.

#### 3.2.9.1 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	

## 3.2.10 PPS IN/OUT Interface

The DPU Controller incorporates an integrated Hardware Clock (PHC) that allows the DPU Controller 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 ASAP<sup>2</sup>). 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.

The DPU Controller PTP solution allows you to run any PTP stack on your host.

With respect to testing and measurements, selected NVIDIA DPU Controllers allow you to use the PPS-out signal from the onboard MMCX RA connector. The DPU Controller also allows 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.

## 3.2.11 External PCIe Power Supply Connector

The external ATX power cable is not supplied with the DPU Controller package; however, this is a standard cable usually available in servers.

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)



### 3.2.12 Cabline CA-II Plus Connectors

The Cabline CA-II connectors on the DPU Controller 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 Controller's component (top) side, whereas the white Cabline CA-II cable mates with the DPU Controller print (bottom) side. The Cabline CA-II cables are offered in three standard lengths; 150mm and 350mm. For connector pinouts, please refer to [Cabline CA-II Plus Connectors Pinouts](#).

### 3.2.13 Integrated BMC Interface

The DPU Controller 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 Controller's power and enables DPU Controller shutdown and power-up.

### 3.2.14 NVMe SSD Interface

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

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

### 3.2.15 RTC Battery

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

### 3.2.16 eMMC Interface

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 below tables 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	SER_DI	
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	SER_DO		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	SER_CLK		B32	GND	
A33	SER_CAPTURE		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	

Pin #	Signal Name	Description	Pin #	Signal Name	Description
A61	PERN10		B61	GND	
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 below 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



Pin Number	Description
4	Sense1
5	GND
6	Sense0
7	GND
8	GND

### 4.3 NC-SI Management Interface

The below tables list 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.
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 ±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

Pin #	Signal Name	I/O	Signal Description
7	RBT_CRSD_V	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
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

Pin #	Signal Name	I/O	Signal Description
17	PRESENCE_N		Presence of DPU. Baseboard should implement a 200 $\Omega$ 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
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

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
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	CIE_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
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

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
51	PCIE_CPU_CX_0P	Micro coax	38	51
52	GND	GND BAR		52
53	No wire	Micro coax	38	53
54	No wire	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_PRSNT1_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	No wire	Micro coax	38	1
2	No wire	Micro coax	38	2
3	No wire	Micro coax	38	3
4	No wire	Micro coax	38	4
5	No wire	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
10	PCIE_CPU_CX_0P	Micro coax	38	10
11	PCIE_CPU_CX_0N	Micro coax	38	11
12	GND	GND BAR		12
13	PCIE_CPU_CX_1P	Micro coax	38	13
14	PCIE_CPU_CX_1N	Micro coax	38	14
15	GND	GND BAR		15
16	PCIE_CPU_CX_2P	Micro coax	38	16
17	PCIE_CPU_CX_2N	Micro coax	38	17
18	GND	GND BAR		18
19	PCIE_CPU_CX_3P	Micro coax	38	19
20	PCIE_CPU_CX_3N	Micro coax	38	20
21	GND	GND BAR		21

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
22	PCIE_CPU_CX_4P	Micro coax	38	22
23	PCIE_CPU_CX_4N	Micro coax	38	23
24	GND	GND BAR		24
25	PCIE_CPU_CX_5P	Micro coax	38	25
26	PCIE_CPU_CX_5N	Micro coax	38	26
27	GND	GND BAR		27
28	PCIE_CPU_CX_6P	Micro coax	38	28
29	PCIE_CPU_CX_6N	Micro coax	38	29
30	GND	GND BAR		30
31	PCIE_CPU_CX_7P	Micro coax	38	31
32	PCIE_CPU_CX_7N	Micro coax	38	32
33	GND	GND BAR		33
34	PCIE_CPU_CX_8P	Micro coax	38	34
35	PCIE_CPU_CX_8N	Micro coax	38	35
36	GND	GND BAR		36
37	PCIE_CPU_CX_9P	Micro coax	38	37
38	PCIE_CPU_CX_9N	Micro coax	38	38
39	GND	GND BAR		39
40	PCIE_CPU_CX_10P	Micro coax	38	40
41	PCIE_CPU_CX_10N	Micro coax	38	41
42	GND	GND BAR		42
43	PCIE_CPU_CX_11P	Micro coax	38	43
44	PCIE_CPU_CX_11N	Micro coax	38	44
45	GND	GND BAR		45
46	PCIE_CPU_CX_12P	Micro coax	38	46
47	PCIE_CPU_CX_12N	Micro coax	38	47
48	GND	GND BAR		48
49	PCIE_CPU_CX_13P	Micro coax	38	49
50	PCIE_CPU_CX_13N	Micro coax	38	50
51	GND	GND BAR		51
52	PCIE_CPU_CX_14P	Micro coax	38	52
53	PCIE_CPU_CX_14N	Micro coax	38	53
54	GND	GND BAR		54
55	PCIE_CPU_CX_15P	Micro coax	38	55
56	PCIE_CPU_CX_15N	Micro coax	38	56
57	GND	GND BAR		57

Pin#	Signal Name	Wire Type	AWG#	Pin# on other end
58	S_PERST1_CONN_L	Micro coax	38	58
59	No wire	No Wire		59
60	S_PRSNT2_L	Micro coax	38	60

# 5 Hardware Installation and PCIe Bifurcation










Installation and initialization of the DPU Controller 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.

	<b>General Installation Instructions</b> Read all installation instructions before connecting the equipment to the power source.
	<b>Jewelry Removal Warning</b> 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.
	<b>Over-temperature</b> 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.
	<b>During Lightning - Electrical Hazard</b> During periods of lightning activity, do not work on the equipment or connect or disconnect cables.
	<b>Copper Cable Connecting/Disconnecting</b> 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.
	<b>Equipment Installation</b> This equipment should be installed, replaced, or serviced only by trained and qualified personnel.
	<b>Equipment Disposal</b> The disposal of this equipment should be in accordance to all national laws and regulations.
	<b>Local and National Electrical Codes</b> This equipment should be installed in compliance with local and national electrical codes.
	<b>Hazardous Radiation Exposure</b> <ul style="list-style-type: none"><li>• 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.</li><li>• 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</li></ul>



## 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 <a href="#">System Requirements</a>
2	Pay attention to the airflow consideration within the host system	Refer to <a href="#">Airflow Requirements</a>
3	Follow the safety precautions	Refer to <a href="#">Safety</a>
4	Unpack the package	Refer to <a href="#">Unpacking</a>
5	Follow the pre-installation checklist	Refer to <a href="#">Pre-Installation Checklist</a>
7	Install the DPU according to the form-factor you have purchased.	Refer to <a href="#">DPU Installation</a>
8	Connect cables or modules to the DPU	Refer to <a href="#">Cables and Modules</a>
9	Power-up the DPU	Refer to <a href="#">DPU Power-Up Instructions</a>

## 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.

Power Supply Requirement
Require a supplementary 8-pin ATX power supply connectivity available through the external power supply connector.

### 5.3.2 Airflow Requirements

The DPU Controller 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 Controller 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 Controller 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 Controllers 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 Controller 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 Controller in a system. Choose the installation instructions according to the configuration you would like to use.

### Installation Instructions

[PCIe x16 Installation Instructions](#)

## Installation Instructions

[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 Controller. 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 Controller before inserting the connector. Do not try and insert the connector upside down. This may damage the DPU Controller.
  - c. Insert the connector into the DPU Controller . 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 Controller.
  - 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.

The BlueField-3 DPU Controller 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 Controller 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 Controller 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 Controller. 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 Controller, 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 Controller 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.

### 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 Controller and the connector lock is secured.

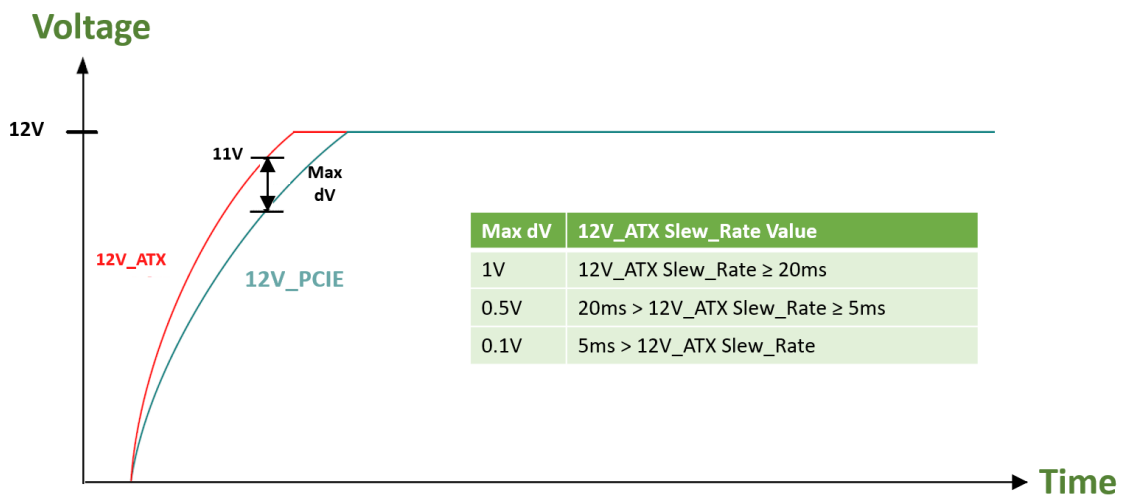
Mechanical Keys of the 8-pin ATX Power Cable (PCIe CEM Specification Rev 5.0)



## 5.9 DPU Controller Power-Up and Power-Down Sequences


### 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.

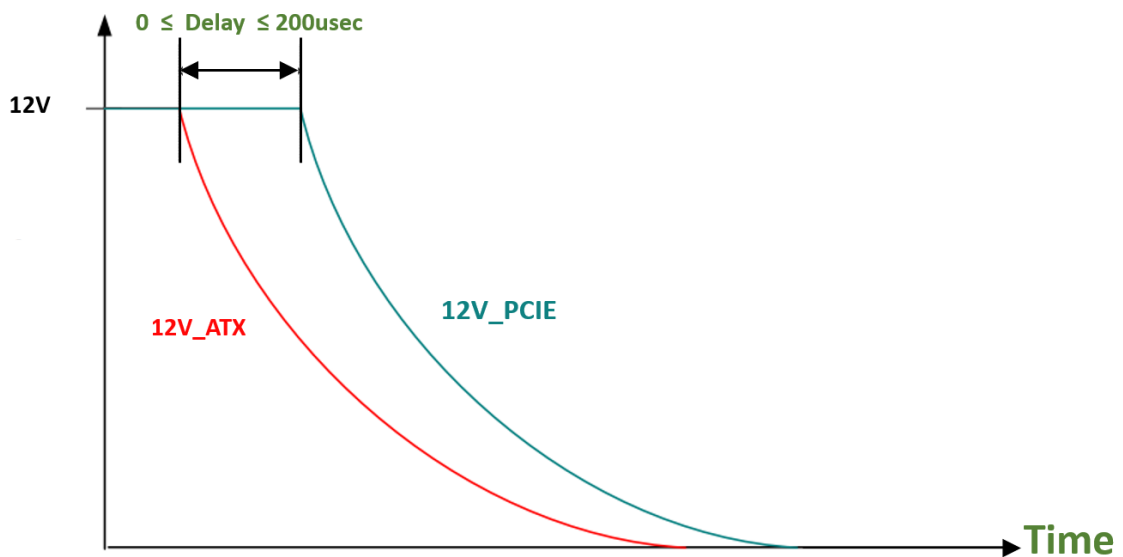
### Voltage



## 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.

### Voltage



2. The 12V\_PCIE voltage must not be powered down while the 12V\_ATX voltage is powered up.


## 5.10 PCIe x16 DPU Controller 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 relevant when the DPU controller is used in Socket Direct mode in dual-socket servers:

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 "WHITE" label should be connected to the connector on the DPU and Auxiliary PCIe card engraved with “White Cable” while the one marked with "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 capable of supplying the required power and airflow as stated in [Specifications](#).

Connect the DPU Controller 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 Controller 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 Controller 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 Controller 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

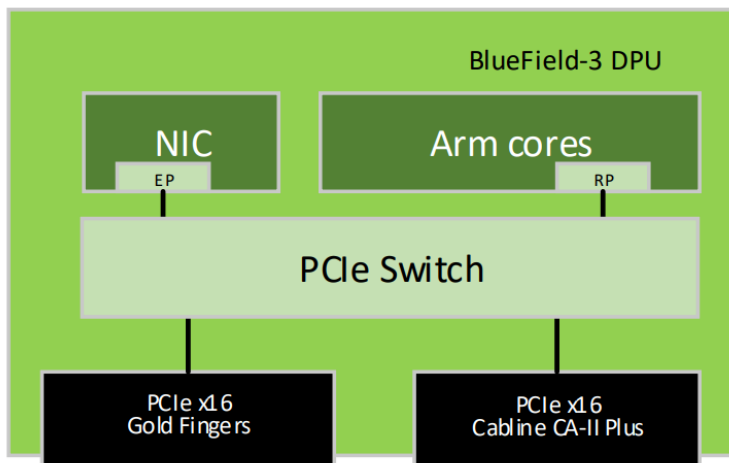
Please note that the following images are for illustration purposes only.

1. Verify that the system is powered off and unplugged.
2. Wait 30 seconds.
3. To remove the card, disengage the retention 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).

NVIDIA BlueField-3 Self-Hosted DPU Controller provides 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, such as 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.



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

- Any customer-set configuration overrides the previous configuration values.
- **WARNING:** 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.

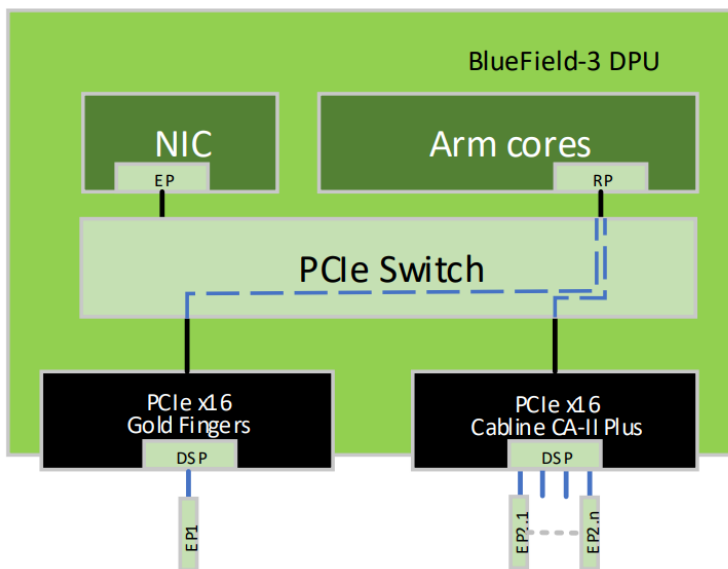
The following table summarizes the available configuration scenarios.

Configuration	Root Port Arm	DPU PCIe Golden Fingers Bifurcation	DPU PCIe Auxiliary Connection Bifurcation
Default	DPU Arm	1 x Gen 5.0/4.0 x16 PCIe lanes as down-stream-port	1 x Gen 5.0/4.0 x16 PCIe lanes as down-stream-port
<a href="#">Bifurcated to Eight x4 PCIe</a>	DPU Arm	4 x Gen 5.0/4.0 x4 PCIe lanes as down-stream-port	4 x Gen 5.0/4.0 x4 PCIe lanes as down-stream-port
<a href="#">Bifurcated to One x16 PCIe and Eight x2 PCIe</a>	DPU Arm	1 x Gen 5.0/4.0 x16 PCIe lanes as down-stream-port	8 x Gen 5.0/4.0 x2 PCIe lanes as down-stream-port

### 5.12.1 Bifurcated to Eight x4 PCIe

In this scenario, the x16 PCIe Goldfingers of the BlueField-3 DPU are bifurcated into four PCIe links, where each link comprises x4 PCIe lanes, 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.

As seen in the below visual representation of this configuration, the DPU's Arm cores function as Root Port of all devices connected to both the Gold Fingers and the Cabline CA-II Plus connector.



#### 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=2
mlxconfig -d <device> s PCI_BUS00_WIDTH=3
mlxconfig -d <device> s PCI_BUS00_SPEED=4
mlxconfig -d <device> s PCI_BUS02_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS02_WIDTH=3
mlxconfig -d <device> s PCI_BUS02_SPEED=4

```



```

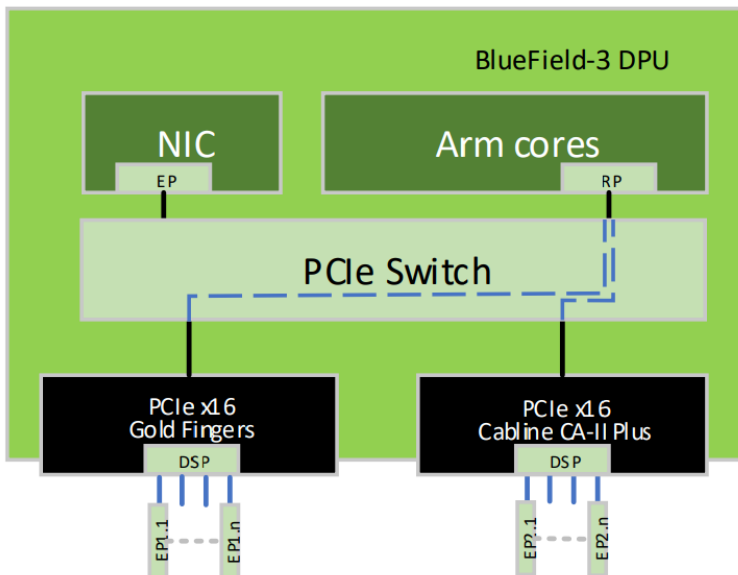
mlxconfig -d <device> s PCI_BUS04_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS04_WIDTH=3
mlxconfig -d <device> s PCI_BUS04_SPEED=4
mlxconfig -d <device> s PCI_BUS06_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS06_WIDTH=3
mlxconfig -d <device> s PCI_BUS06_SPEED=4
mlxconfig -d <device> s PCI_BUS10_HIERARCHY_TYPE=2
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

```

## 5.12.2 Bifurcated to One x16 PCIe and Eight x2 PCIe

In this scenario, the x16 PCIe Goldfingers of the BlueField-3 DPU are bifurcated to one PCIe link with x16 PCIe lanes, while the additional x16 PCIe lanes are accessible via the Cabline CA-II Plus connector, bifurcated into eight PCIe links, where each link comprises x2 PCIe lanes.

As seen in the below visual representation of this configuration, the DPU's Arm cores function as Root Port of all downstream devices connected to both the Goldfingers and the Cabline CA-II Plus connector.



### 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=2
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=2
mlxconfig -d <device> s PCI_BUS10_SPEED=4
mlxconfig -d <device> s PCI_BUS11_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS11_WIDTH=2
mlxconfig -d <device> s PCI_BUS11_SPEED=4

```

```
mlxconfig -d <device> s PCI_BUS12_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS12_WIDTH=2
mlxconfig -d <device> s PCI_BUS12_SPEED=4
mlxconfig -d <device> s PCI_BUS13_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS13_WIDTH=2
mlxconfig -d <device> s PCI_BUS13_SPEED=4
mlxconfig -d <device> s PCI_BUS14_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS14_WIDTH=2
mlxconfig -d <device> s PCI_BUS14_SPEED=4
mlxconfig -d <device> s PCI_BUS15_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS15_WIDTH=2
mlxconfig -d <device> s PCI_BUS15_SPEED=4
mlxconfig -d <device> s PCI_BUS16_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS16_WIDTH=2
mlxconfig -d <device> s PCI_BUS16_SPEED=4
mlxconfig -d <device> s PCI_BUS17_HIERARCHY_TYPE=2
mlxconfig -d <device> s PCI_BUS17_WIDTH=2
mlxconfig -d <device> s PCI_BUS17_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.

OPN	Data Transmission Rate	Default Protocol and Rate
900-9D3C6-00CV-GA0 900-9D3C6-00CV-DA0	InfiniBand: NDR200 200Gb/s Ethernet: 200GbE	Ethernet 200GbE

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

UEFI can configure the DPU 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

## 7.1 General Troubleshooting

<b>Server unable to find the DPU</b>	<ul style="list-style-type: none"><li>• Ensure that the DPU is placed correctly</li><li>• Make sure the DPU slot and the DPU are compatible Install the DPU in a different PCI Express slot</li><li>• Use the drivers that came with the DPU or download the latest</li><li>• Make sure your motherboard has the latest BIOS</li><li>• Try to reboot the server</li></ul>
<b>The DPU no longer works</b>	<ul style="list-style-type: none"><li>• Reseat the DPU in its slot or a different slot, if necessary</li><li>• Try using another cable</li><li>• Reinstall the drivers for the network driver files may be damaged or deleted</li><li>• Reboot the server</li></ul>
<b>DPU's stopped working after installing another DPU</b>	<ul style="list-style-type: none"><li>• Try removing and re-installing all DPUs</li><li>• Check that cables are connected properly</li><li>• Make sure your motherboard has the latest BIOS</li></ul>
<b>Link indicator light is off</b>	<ul style="list-style-type: none"><li>• Try another port on the switch</li><li>• Make sure the cable is securely attached</li><li>• Check you are using the proper cables that do not exceed the recommended lengths</li><li>• Verify that your switch and DPU port are compatible</li></ul>
<b>Link light is on, but with no communication established</b>	<ul style="list-style-type: none"><li>• Check that the latest driver is loaded</li><li>• Check that the DPU and its link are set to the same speed, and duplex settings</li></ul>
<b>Forgot password needed to install/upgrade the DPU image</b>	Refer to the latest version of <a href="#">BlueField DPU SW Manual</a> and follow the instructions under "Upgrading NVIDIA BlueField DPU Software" section.

## 8 Specifications

The DPU's maximum power consumption does not exceed 150W 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

### 8.1 900-9D3C6-00CV-GA0 / 900-9D3C6-00CV-DA0 Specifications


<b>BlueField-3 IC</b>	BlueField-3 E-Series - 16 Arm-Cores - 505MHz/2000MHz		
	Crypto Enabled with integrated BMC		
<b>Physical</b>	Single-Slot FHHL Card Dimensions: 111.15mm x 167.65mm Tall Bracket Dimensions: 121.0mm x 21.6mm 900-9D3C6-00CV-DA0 Only: Heatsink Dimensions (Length, Width, Height): 139.6mm x 92.7mm x 10.2mm		
<b>Interfaces</b>	See <a href="#">Supported Interfaces</a>		
	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 5.0 @ SERDES 32GT/s through the PCIe auxiliary passive card and Cabline SA-II Plus harnesses	
	On-Board Memory	<ul style="list-style-type: none"> <li>• Dual-channel with 20 DDR5 + ECC (128bit + 16bit ECC) for a total of 48GB @ 5200MT/s</li> <li>• 40GB pSLC eMMC memory + 128GB SSD</li> </ul>	
Networking Connector	Dual-port QSFP112 (copper and optical)		
<b>Data Rate</b>	InfiniBand	NDR200/HDR/HDR100/EDR/FDR/SDR	
	Ethernet (Default Speed)	200/100/50/25/10 Gb/s	
<b>Protocol Support</b>	InfiniBand: IBTA v1.5a 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, 200GAUI-2 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, XLPII, SFI		
<b>Electrical and Thermal Specifications</b>	Voltage: 12V		
	Electrical and thermal specifications are provided in "NVIDIA BlueField-3 DPUs and SuperNICs Electrical and Thermal Specifications" document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.		
<b>Environmental</b>	Temperature	Operational	0 °C to 55 °C
		Non-operational	-40 °C to 70 °C <sup>a</sup>
	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	


a. The non-operational storage temperature specifications apply to the product without its package.

### 8.1.1 DPU Controller 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.



#### 8.1.1.1 Tall Bracket Mechanical Drawing and Dimensions



---

## 9 Monitoring

### 9.1 Thermal Sensors

The thermal threshold definitions are described in the "NVIDIA BlueField-3 Networking Platforms Electrical and Thermal Specifications" document. You can access the document either by logging into NVOnline or by contacting your NVIDIA representative.

### 9.2 Heatsink

This section applies to 900-9D3C6-00CV-DA0 only, 900-9D3C6-00CV-GA0 is shipped without a heatsink.

The heatsink is attached to the DPU by three screws to dissipate the heat from the IC. The DPU IC 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 below table for heatsink details per card configuration. For the required airflow (LFM) per OPN, please refer to [Specifications](#).

Maximum Dimensions
Length, Width, Height: 139.6mm x 92.7mm x 10.2mm

# 10 Finding the GUID/MAC and Serial Number on the Card

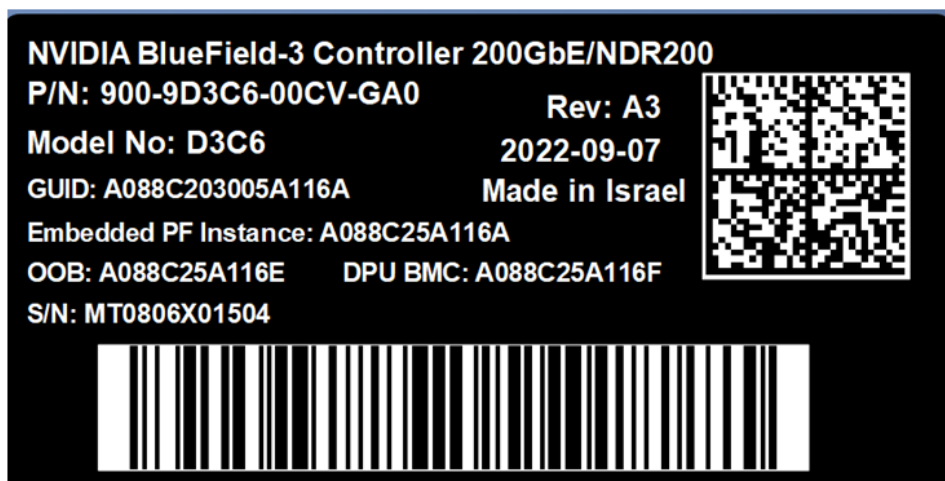
Each BlueField-3 platform is uniquely identified to facilitate efficient management and configuration. This section aims to guide you through identifying your BlueField-3 platform by using the information provided on the board label.

The board label includes unique identifiers essential for network communication—specifically, a primary GUID address for InfiniBand protocols or 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.

Refer to the board label example and a table listing the different MAC addresses, offering examples and additional information to enhance clarity.

## 10.1 DPU Controller 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)	GUIDs is a unique 16 byte identifier of the BlueField in an InfiniBand network, the GUID never changes.	The GUID on the label is assigned to the <b>System GUID</b> 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 <b>Node GUID</b> of <b>port 0</b> is equal to the System GUID. The <b>Node GUID</b> of <b>port 1</b> 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	<p>DPU's with more than one high-speed port will use incremental (+1) MAC addresses against the base HOST MAC on the label.</p> <p>For instance:  From the label - HOST: 94 6D AE F5 A1 CC  <b>Port 0:</b> 94 6D AE F5 A1 CC (same as label)  <b>Port 1:</b> 94 6D AE F5 A1 CD (+1 from base)  <b>Port 2:</b> 94 6D AE F5 A1 CE (+2 from base)  <b>Port 3:</b> 94 6D AE F5 A1 CF (+3 from base)</p>
OOB (Out-of-Band)	OOB represents the MAC address of the 1GbE out-of-band management port of 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.
DPU BMC (Baseboard Management Controller)	Represents the MAC address for the 1GbE OOB management port of the DPU BMC. Note: The base MAC address for DPU BMC is +1 from the OOB MAC address.	
S/N (Serial Number)	A unique serial number that distinguishes your DPU from others. The S/N is used for tracking and support purposes.	N/A

# 11 PCIe Auxiliary Card Kit

## Down Stream Port (DSP)

The self-hosted BlueField-3 DPU controllers 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. 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.

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

## 11.1 PCIe Auxiliary Card Package Contents

Category	Qty	Item
Cards	1	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
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)

## 11.2 Channel Insertion Loss

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

The following table describes the BlueField-3 DPU 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 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 Golden-Fingers) to be 9.5dB @16Ghz.

The BlueField-3 AIC, together with a 15cm Cabline CA-II Plus harnesses and the Passive PCIe Auxiliary Card loss is: 3dB+2dB+3.8dB+1.5dB=10.3dB > 9.5dB

## 11.3 Cabline CA-II Plus Harness Pinouts

### 11.3.1 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 Network Adapter 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 Network Adapter 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	No wire	Micro coax		38	53
54	No wire	Micro coax		38	54
55	Reserved_01	Micro coax		38	55
56	No wire	Micro coax		38	56
57	I2C_AUX_SCL	Micro coax	The ConnectX-7 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 network adapter 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 network adapter 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 ConnectX-7 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 network adapter 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 network adapter 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

### 11.3.2 Cabline CA-II Plus Harness - Print Side

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
1	No wire	Micro coax	38		1
2	No wire	Micro coax	38		2
3	No wire	Micro coax	38		3
4	No wire	Micro coax	38		4

Pin#	Signal Name	Wire Type	AWG#	Detailed Description	Pin# on the other end
5	No wire	Micro coax	38		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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 ConnectX-7 card.	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 network adapter. When connecting a GPU or an end point to the PCIe Cabline CA-II Plus interface, this signal is driven from the network adapter side (which operates as a PCIe switch in this case), to the GPU or end-point on the motherboard side.	58
59	No wire	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 network adapter main card. 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


## 11.4 Technical Specifications

<b>Physical</b>	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		
<b>Power Consumption</b>	<b>Voltage:</b> 12V, 3.3V_AUX <b>Maximum current:</b> 100mA for the 3.3V_AUX voltage rail		
<b>PCIe Connectivity</b>	930-9DAX5-0015-000	PCI Express Gen 5.0/4.0: SERDES @ 16/32 GT/s, x16 lanes (Gen 3.0 compatible)	
	930-9DAX4-0035-000	PCI Express Gen 4.0: SERDES @ 16GT/s, x16 lanes (Gen 3.0 compatible)	
<b>Environmental</b>	<b>Temperature</b>	<b>Operational</b>	0° C to 55° C
		<b>Non-operational</b>	-40° C to 70° C
	<b>Humidity</b>	<b>Operational</b>	10% to 85% relative humidity
		<b>Non-operational</b>	10% to 90% relative humidity
<b>Altitude (Operational)</b>	3050m		
<b>Regulatory</b>	<b>Safety:</b> CB / cTUVus / CE		
	<b>EMC:</b> CE / FCC / VCCI / ICES / RCM / KC		

<b>Physical</b>	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
<b>Power Consumption</b>	<b>Voltage:</b> 12V, 3.3V_AUX <b>Maximum current:</b> 100mA for the 3.3V_AUX voltage rail
	<b>RoHS:</b> RoHS Compliant

### 11.4.1 PCIe Auxiliary Card Mechanical Drawings and Dimensions


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



### 11.4.2 Bracket Mechanical Drawings and Dimensions

Auxiliary PCIe Connection Card Tall Bracket	Auxiliary PCIe Connection Card Short Bracket
<p>The tall bracket has a length of 121.05 mm and a height of 21.59 mm. It features a series of mounting holes along its length.</p>	<p>The short bracket has a length of 80.29 mm and a height of 22.83 mm. It also features a series of mounting holes along its length.</p>

11.4.3 Cabline CA-II Plus Harnesses Mechanical Drawing



---

## 12 Cables and Cabling Configurations

### 12.1 RJ45 Cable for the 1GbE OOB Management Interface



### 12.2 USB Interface Cabling

Debugging and loading new versions of the operating systems and firmware requires the use of Type A to Type B USB 2.0 Cable. The following figure shows an example of Type A to Type V USB Cable.




# 13 General Cooling Recommendations for Heatsink-less DPU Controllers

This section offers general cooling recommendations to NVIDIA customers who have acquired a heatsink-less DPU Controller (OPN: 900-9D3C6-00CV-GA0).

### IMPORTANT NOTES:

1. The customer holds exclusive responsibility for the thermal design and for ensuring all board components never exceed their designated thermal operating limits.
2. The recommendations provided in this section are based on an air-cooled thermal design.
3. If the customer's environment has low or no airflow (e.g., liquid cooling) where the DPU Controller is to be installed, the customer must ensure proper cooling for all the other board components that are not covered by the recommendations of this chapter or in the product's thermal model. The customer must guarantee those components do not exceed their thermal operating limits as well.

Isometric view of the heatsink-less DPU Controller:



## 13.1 General Notes

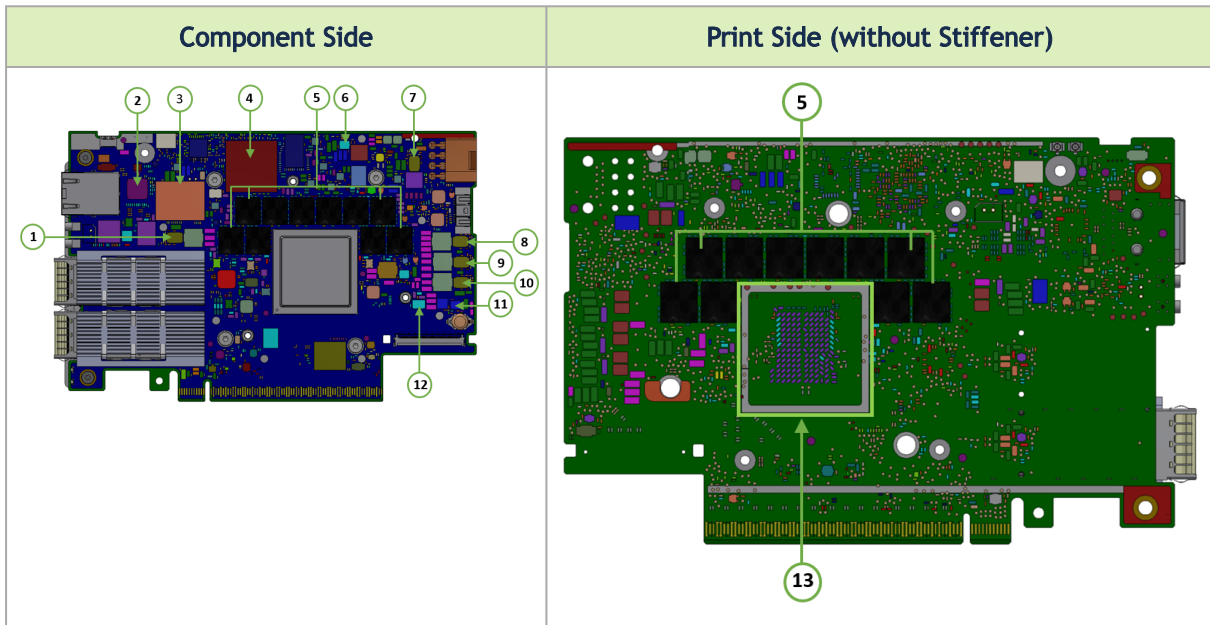
1. The Thermal Transfer Plate (TTP) may encompass an internal liquid cooling system within a cold plate.
2. The M3 screws maximum protrusion length according to stiffener thread depth is listed in [Mechanical Interfaces](#).
3. Recommended torque values: M3: 0.35-0.4 Nm. Use a locking mechanism.
4. To ensure unobstructed access for replacement, ensure that the battery is not blocked. Refer to [Encapsulation and Battery on SFG](#).

- When considering encapsulation, ensure there is no interference with the TTP. The encapsulation is presented as the maximum tolerance on the SFG and is labeled with ENC. Refer to [Encapsulation and Battery on SFG](#).

## 13.2 Components that Require Cooling

It is the responsibility of the customer to ensure proper cooling and thermal regulation for all components on the board. The provided guidelines are minimal and are based on NVIDIA's air-cooling thermal solution. Note the following:

- The DPU controller SoC require contact with TTP on the component side of the product.
- Within the air-cooled thermal solution, the cooling of components on the print side of the DPU controller is achieved via attachment to the board's bottom stiffener with the thermal interface material. See example in the board print side isometric view above.
- The battery's thermal limitation is the maximum ambient temperature of 60°C.
- The acceptable pressure on the SoC adheres to the following specifications:
  - 35 psi (SoC area 761.5 mm<sup>2</sup>).
  - The center force should be applied at the center of the SoC.



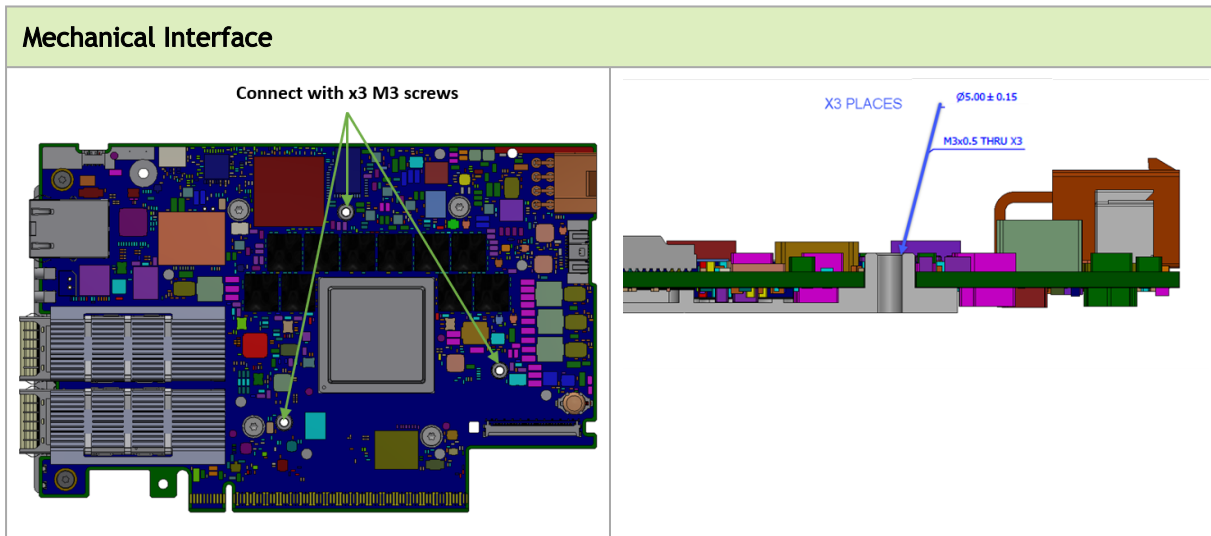
Item #	Board Components
1	U63
2	U44
3	M1394
4	U99
5	M111-M115, M221-M225, M331-M335, M441-M445
6	U30
7	U66
8	U62



Item #	Board Components
9	U64
10	U65
11	U32
12	U33
13	See marked square


### 13.2.1 Mechanical Interfaces

The following figure lists the mechanical interfaces of the DPU controller.



### 13.2.2 Encapsulation and Battery on SFG

Ensure that the battery remains unobstructed to allow access for replacement. Furthermore, avoid any interference with the TTP. Note that the encapsulation is depicted as the maximum tolerance in the SFG and is labeled as ENC.



---

## 14 Document Revision History

Date	Comments/Changes
Apr. 2024	<ul style="list-style-type: none"><li>Removed the "Securing Workloads" feature from <a href="#">Features and Benefits</a></li><li>Updated the important notes and warnings in <a href="#">External PCIe Power Supply Connector</a></li><li>Added <a href="#">8-pin ATX Power Supply Cable</a> installation instructions</li><li>Updated <a href="#">Cabline CA-II Plus Harness Pinouts</a></li><li>Updated <a href="#">PCIe Auxiliary Card Kit</a></li><li>Added notes across the document stating that the Socket-Direct functionality is currently not supported by firmware</li></ul>
Nov. 2023	<ul style="list-style-type: none"><li>Updated <a href="#">PCIe Auxiliary Card Kit</a></li></ul>
Oct. 2023	<ul style="list-style-type: none"><li>Updated <a href="#">NVMe SSD Interface</a></li><li>Updated <a href="#">DPU Power-Up and Power-Down Instructions</a></li></ul>
Aug. 2023	<ul style="list-style-type: none"><li>Added <a href="#">General Cooling Recommendations for Heatsink-less DPU Controllers</a></li><li>Updated <a href="#">Cabline CA-II Plus Connectors Pinouts</a>.</li><li>Updated <a href="#">Cabline CA-II Plus Harness Pinouts</a>.</li><li>Added a note to <a href="#">NVMe SSD Interface</a>.</li></ul>
May. 2023	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.

