NVIDIA BlueField-2 BF2500 Ethernet DPU Controller User Manual

Exported on Oct/31/2023 01:37 PM

Table of Contents

About This Manual

It provides details as to the interfaces of the board, specifications, required software and firmware for operating the board, and a step-by-step plan of how to bring up the BlueField-2 DPU.

EOL'ed (End of Life) Ordering Part Numbers

OPN	FormFactor	Series/ Core Speed	Max Speed	No. of Ports	PCIe Support	Crypto	Secure Boot	1GbE OOB	On-board DDR Memory	On-board eMMC Memory
MBF2H516B-CEEOT	FHHL*	P-Series / 2.75GHz	100GbE	2xQSFP56	Gen 4.0 x16	Crypto Enabled	-	✓	16GB	64GB
MBF2H516B-CENOT	FHHL	P-Series / 2.75GHz	100GbE	2xQSFP56	Gen 4.0 x16	Crypto Disabled	-	✓	16GB	64GB

* Full-Height Half-Length

Intended Audience

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

Technical Support

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

- E-mail: Enterprisesupport@nvidia.com
- Enterprise Support page: https://www.nvidia.com/en-us/support/enterprise

Related Documentation

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

Document Conventions

When discussing memory sizes, GB and GBytes are used in this document to mean size in 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.

Introduction

This is the User Manual for NVIDIA® BlueField®-2 BF2500 DPU Controller. This document provides details of the product interfaces, specifications, required software and firmware for operating the board, and a step-by-step plan of how to bring up the BlueField-2 BF2500 DPU Controller.

System Requirements

Item	Description
Main-board PCI Express slot	x16 Gen 4.0 slot.
System Power Supply	Minimum 75W or greater system power supply for all cards.
	These PCIe Gen 4.0 x16 DPU controllers require additional 75W through a supplementary 6-pin ATX power supply connector. **NOTE:** The connector is not included in the package. It should be part of system wiring or it can be ordered separately as a system accessory.
Operating System	BlueField-2 DPU 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. BlueField-2 DPU also supports CentOS and has an out-of-band 1GbE management interface. For more information, please refer to the DOCA SDK documentation or NVIDIA BlueField-2 Software User Manual.
Connectivity	• Interoperable with 1/10/25/40/50/100/200 Gb/s Ethernet switches • Passive copper cable with ESD protection • Powered connectors for optical and active cable support

For detailed information, see Specifications.

Package Contents

Before installing your new system, unpack it and check against the below tables that all the parts have been sent. Check the parts for visible damage that may have occurred during shipping.

> If anything is damaged or missing, contact your reseller.

Card Package

Item	Description
Cards	1x BlueField-2 DPU Controller card with an assembled tall bracket

Accessories Kit

The accessories kit should be ordered separately. Earlier controller versions require the kit OPN MBF20-DKIT, while newer versions require kit OPN MBF25-DKIT.

Kit OPN	Contents
MBF20-DKIT	1x USB 2.0 Type A to mini-USB Type B cable
	1x USB 2.0 Type A to 30pin Flat Socket

> These DPU controllers, you need a 6-pin ATX power supply connector cable to activate the card. The cable is not included in the package. For further details, please refer to External PCIe Power Supply Connector.

Product Overview

BlueField-2 BlueField-2 BF2500 DPU Controller features the second generation BlueField-2 data processing Unit (DPU) – an innovative and high-performance programmable networking engine. The DPU integrates an array of eight powerful 64-bit Arm v8 A72 cores interconnected by a coherent mesh with a DDR4 memory controller and a dual-port Ethernet network controller. Providing unmatched scalability and efficiency, NVIDIA BF2500 DPU Controller is the ideal adapter to accelerate the most demanding workloads in data center, cloud, service provider and storage environments.

⚠ The BlueField BF2500 DPU Controller should be installed only in a JBOF and JBOD Systems as it functions as a PCIe root-complex (RC) initiating PCIe bus operations. Installing it in a regular host system may damage the card.

Ideal Solution for JBOF and JBOD Systems

NVIDIA BlueField-2 DPU is a highly integrated and efficient controller, optimized for NVMe storage systems, Network Functions Virtualization (NFV), Cloud and Machine Learning workloads. BlueField-2 integrates all the discrete components of a storage system appliance into a single chip, including Arm core CPUs, PCIe switch and a network controller, 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 offload accelerator, the BF2500 DPU Controller has a superior performance advantage over existing JBOF systems, significantly reducing storage transaction latency, while increasing IOPs (I/O operations per second).

Features and Benefits

This section describes hardware features and capabilities.

⚠ It is recommended to upgrade your BlueField product to the latest software and firmware versions available in order to enjoy the latest features and bug fixes.
Please refer to the software release notes for feature availability.

Feature	Description
PCI Express (PCIe)	Uses PCIe Gen 4.0 (16GT/s) through an x16 edge connector, compatible with Gen 3.0, 2.0 and 1.1.

Feature	Description
Up to 200 Gigabit Ethernet	• The adapters comply with the following IEEE 802.3 standards: 200Gb/s / 100Gb/s / 50Gb/s / 40Gb/s / 25Gb/s / 10Gb/s / 1Gb/s • IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet • IEEE 802.3by, Ethernet Consortium25, 50 Gigabit Ethernet, supporting all FEC modes • IEEE 802.3ba 40 Gigabit Ethernet • IEEE 802.3by 25 Gigabit Ethernet • IEEE 802.3ae 10 Gigabit Ethernet • IEEE 802.3ap based auto-negotiation and KR startup • IEEE 802.3ad, 802.1AX Link Aggregation • IEEE 802.1Q, 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 • Jumbo frame support (9.6KB)
On-board Memory	• Quad SPI NOR FLASH - includes 256Mbit for Firmware image. • UVPS EEPROM - includes 1Mbit. • 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. • eMMC - x8 NAND flash • DDR4 SDRAM - 16GB @3200MT/s single-channel DDR4 SDRAM memory. Solder down on-board. 64bit + 8bit ECC.
BlueField-2 DPU	The BlueField-2 DPU integrates eight 64-bit Armv8 A72 cores interconnected by a coherent mesh network, one DRAM controller, an RDMA intelligent network adapter supporting up to 200Gb/s, an embedded PCIe switch with endpoint and root complex functionality, and up to 16 lanes of PCIe Gen 4.0.
Overlay Networks	In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. DPU effectively addresses this by providing advanced NVGRE and VXLAN hardware offloading engines that encapsulate and de-capsulate the overlay protocol.
RDMA and RDMA over Converged Ethernet (RoCE)	DPU, utilizing IBTA RDMA (Remote Data Memory Access) and RoCE (RDMA over Converged Ethernet) technology, delivers low-latency and high-performance over 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.
NVIDIA PeerDirect	NVIDIA PeerDirect communication provides high-efficiency RDMA access by eliminating unnecessary internal data copies between components on the PCIe bus (for example, from GPU to CPU), and therefore significantly reduces application run time. DPU advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.

Feature	Description
Quality of Service (QoS)	Support for port-based Quality of Service enabling various application requirements for latency and SLA.
Storage Acceleration	A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage RDMA for high-performance storage access. • NVMe over Fabric offloads for the target machine • T10-DIF Signature Handover BlueField-2 DPU 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 BlueField-2 Arm cores. On the storage initiator side, BlueField-2 DPU can prove an efficient solution for hyper-converged systems to enable the host CPU to focus on compute while all the storage interface is handled through the Arm cores.
NVMe-oF	Nonvolatile Memory Express (NVMe) over Fabrics is a protocol for communicating block storage IO requests over RDMA to transfer data between a host computer and a target solid-state storage device or system over a network. BlueField-2 DPU may operate as a co-processor offloading specific storage tasks from the host using its powerful NVMe over Fabrics Offload accelerator.
SR-IOV	DPU SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server.
GPU Direct	The latest advancement in GPU-GPU communications is GPUDirect RDMA. This new technology provides a direct P2P (Peer-to-Peer) data path between the GPU Memory directly to/from the 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. DPU uses high-speed DMA transfers to copy data between P2P devices resulting in more efficient system applications
Crypto	The BlueField-2 DPU crypto enabled versions include a BlueField-2 IC which supports accelerated cryptographic operations. In addition to specialized instructions for bulk cryptographic processing in the Arm cores, an offload hardware engine accelerates public-key cryptography and random number generation are enabled.
Security Accelerators	A consolidated compute and network solution based on DPU achieves significant advantages over a centralized security server solution. Standard encryption protocols and security applications can leverage BlueField-2 compute capabilities and network offloads for security application solutions such as Layer4 Stateful Firewall.
Out-of-Band Management	The BlueField-2 DPU 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.

Supported Interfaces

Layout and Interfaces

This section describes the DPU supported interfaces. Each numbered interface that is 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 BF2500 card.

Component Side

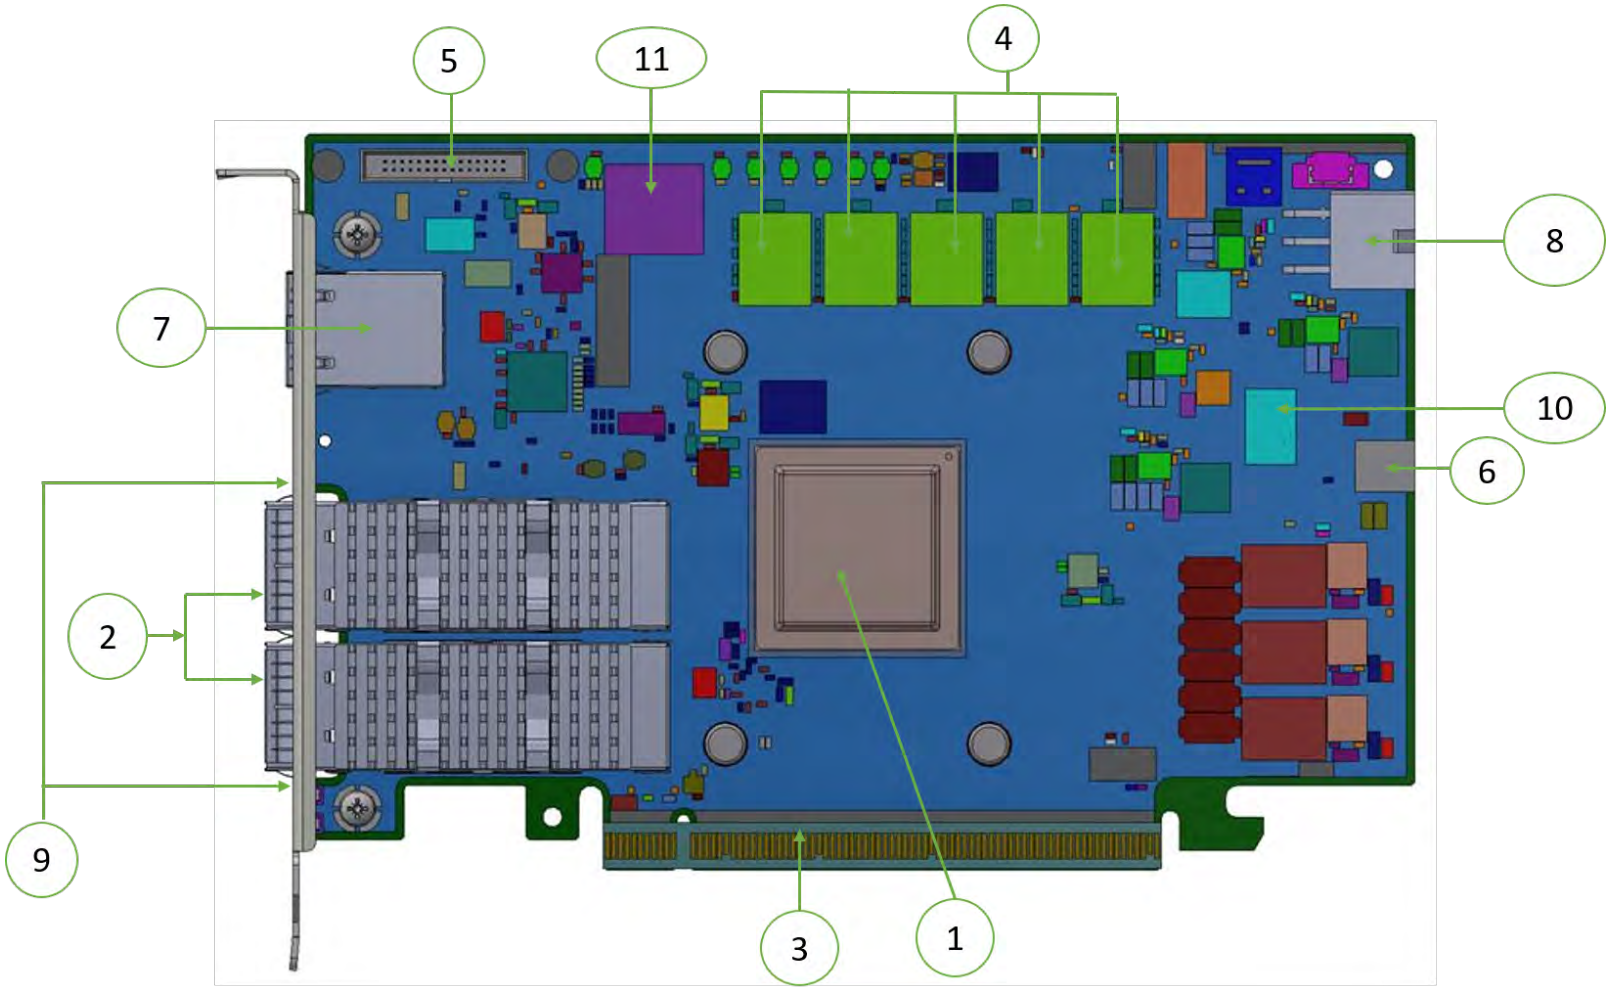

Print Side

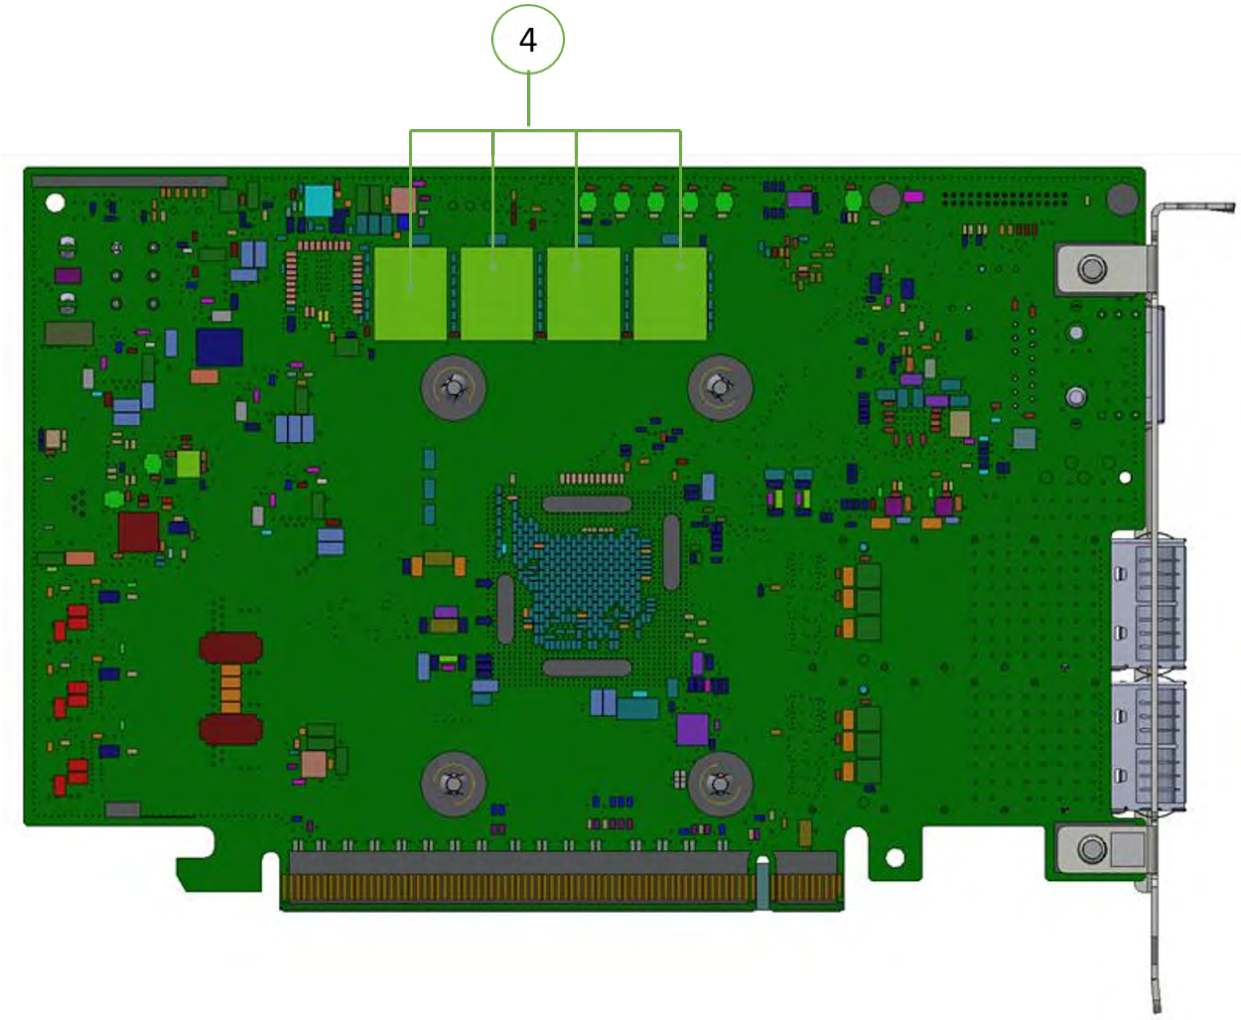
4

Item	Interface	Description
1	DPU	DPU IC 8 cores
2	Ethernet QSFP56 Interface	The ethernet traffic is transmitted through the DPU QSFP56 connectors. The QSFP56 connectors allow the use of modules, optical and passive cable interconnect solutions. By default, the ports of this group of OPNs are set to operate in QSFP28 mode (default card firmware setting).
3	PCI Express Interface	PCIe Gen 4.0 through an x16 edge connector
4	DDR4 SDRAM On-Board Memory	Units of SDRAM for a total of 16 GB @ 3200MT/s single DDR4 channel, 64bit + 8bit ECC, solder-down memory
5	NC-SI Management Interface	BMC connectivity for remote management
6	Mini USB Type B Interface	Used for OS image loading
7	1GbE OOB Management Interface	1GbE BASE-T OOB management interface
8	External PCIe Power Supply Connector	An external 12V power connection through a 6-pin ATX connector.
9	Networking Ports LEDs Interface	One bi-color I/O LED per port to indicate link and physical status
10	RTC Battery	CR621 battery holder for RTC
11	eMMC Interface	x8 NAND flash

Interfaces Detailed Description

DPU

NVIDIA® BlueField®-2 DPU is a family of advanced DPU IC solutions that integrate a coherent mesh of 64-bit Arm v8 A72 cores, an NVIDIA® ConnectX®-6 Dx network adapter front-end and a PCI Express switch into a single chip. The powerful DPU IC architecture includes an Arm v8 multicore processor array and enables 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 BlueField-2, the ConnectX-6 Dx 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.

Encryption

⚠ Applies to Crypto enabled OPNs.

DPU 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 individual 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 is capable of delivering powerful functionality, including encryption of data-in-motion, bare-metal provisioning, stateful L4 firewall and more.

Ethernet QSFP56 Interface

The network ports of the DPU are compliant with the IEEE 802.3 Ethernet standards listed in Features and Benefits. Ethernet traffic is transmitted through the cards' QSFP56 connectors. Note that the ports operate in QSFP28 mode by default.

PCI Express Interface

The DPU supports PCI Express Gen4.0 (3.0, 2.0, and 1.1 compatible) through an x16 edge connector. The following lists PCIe interface features:

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

DDR4 SDRAM On-Board Memory

The DPU incorporates 16GB @ 3200MT/s single DDR4 channel, 64bit + 8bit ECC, solder-down memory.

NC-SI Management Interface

The DPU enables the connection of a Baseboard Management Controller (BMC) to a set of Network Interface Controller (NICs) for the purpose of enabling out-of-band remote manageability. The NC-SI management is supported over RMII and has a connector on the DPU. Please refer to NC-SI Management Interface for pins.

UART Interface Connectivity

A UART debug interface is available on the DPU cards via 3 pins of a 30-pin NC-SI connector (described in NC-SI Management Interface). The connectivity is shown in the following table:

NC-SI Connector Pin #	Signal on DPU
30	BF_UART0_RX
28	BF_UART0_TX
26	GND

The UART interface is compliant with TTL 3.3V voltage level. A USB to UART cable that supports TTL voltage levels should be used to connect the UART Interface for Arm console access - see example below.

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

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

USB Interfaces

The controllers use a mini-USB Type B connector to load operating system images.

1GbE OOB Management Interface

The DPU 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.

10Mb/s and 100Mb/s modes are not supported on this interface.

1GbE OOB Management LEDs Interface

There are 2 OOB management LEDs, one green and one amber/yellow. The following table describes LED behavior for DPUs with or with on-board BMC.

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

RTC Battery

The DPU incorporates a COIN TYPE LITHIUM BATTERY CR621 for RTC (Real Time Clock).

eMMC Interface

The DPU incorporates an eMMC interface on the card's print side. The eMMC is an x8 NAND flash and is used for Arm boot, operating system storage and disk space. Memory size is 64GB.

External PCIe Power Supply Connector

Applies to FHHL P-Series DPUs with x16 PCIe Gen 4 lanes only, which require supplementary power to be fed via the DPU's on-board 6-pin ATX power supply connector. The power cable that should be connected to this on-board ATX connector is *not* supplied with the DPU; however, this is a standard cable that is normally available in servers.

The FHHL P-Series DPUs with x16 PCIe Gen 4 lanes incorporate an external 12V power connection through a 6-pin ATX connector. The DPU includes a special circuitry that provides current balancing between the two power supplies; the 12V from the PCIe x16 standard slot and the 12V from the ATX 6-pin connector. Since the power provided by the PCIe golden fingers is limited to 75W, a total maximum of up to 150W is enabled through both the ATX 6-pin connector and the PCIe x16 golden fingers. The actual power consumption is in accordance to the mode of operation of the DPU and is split evenly between the two power sources.

For the pinout of the on-board 6-pin ATX connector, please refer to External PCIe Power Supply Connector Pins.

Networking Ports LEDs Interface

There is one bicolor (Yellow and Green) I/O LED per port to indicate speed and link status.

Link Indications

<table>
<tr><th>State</th><th>Bi-Color LED (Yellow/Green)</th></tr>
<tr><td>Beacon command for locating the adapter card</td><td>1Hz blinking Yellow</td></tr>
<tr><td>Error</td><td>4Hz blinking Yellow Indicates an error with the link. The error can be one of the following:
<table>
<tr><th>Error Type</th><th>Description</th><th>LED Behavior</th></tr>
<tr><td>I^2C</td><td>I^2C access to the networking ports fails</td><td>Blinks until error is fixed</td></tr>
<tr><td>Over-current</td><td>Over-current condition of the networking ports</td><td>Blinks until error is fixed</td></tr>
</table>
</td></tr>
<tr><td>Physical Activity</td><td>• A constant Green indicates a link with the maximum networking speed.
• A constant Yellow indicates a link with less than the maximum networking speed.</td></tr>
<tr><td>Link Up</td><td>• A constant Green indicates a link with the maximum networking speed.
• A constant Yellow indicates a link with less than the maximum networking speed.</td></tr>
</table>

Pin Description

PCI Express Interface

The BF2500 Controller Card uses a PCI Express x16 connector (component and print sides) according to the PCI Express 4.0 specifications. The following two tables provide component side pins description and print side pins description.

Pin #	Signal Name	I/O	Description	Pin #	Signal Name	I/O	Description
A1	PRSNT1#	Input	Card presence detect, connected to PRSNT2#	B1	12V	Power	12V
A2	12V	Power	12V	B2	12V	Power	12V
A3	12V	Power	12V	B3	12V	Power	12V
A4	GND	GND	Ground	B4	GND	GND	Ground
A5	TCK		JTAG - Not Connected	B5	SMCLK	I/O	SMBUS_SCL - connected to DPU SoC
A6	TDI	Input	On-board loop	B6	SMDAT	I/O	SMBUS_SDA - connected to DPU SoC
A7	TDO	Output		B7	GND	GND	Ground
A8	TMS		JTAG - Not Connected	B8	3.3V	Power	3.3V - Not Connected
A9	3.3V	Power	3.3V - Not Connected	B9	TRST#		JTAG - Not Connected
A10	3.3V	Power	3.3V - Not Connected	B10	3.3V_AUX	Power	3.3V auxiliary power
A11	PERST#	Input	PCIe Interface Reset#, driven by DPU SoC	B11	WAKE#/RSVD		Not connected, option connect to NIC_GPIO[13]

A12	GND	GND	Ground	**B12**	RSVD		Not connected
A13	REFCLK+	Output	PCIe Reference Differential Clock. Not connected (default)	**B13**	GND	GND	Ground
A14	REFCLK-	Output	PCIe Reference Differential Clock. Not connected (default)	**B14**	PETP0	Input	DPU SoC RX
A15	GND	GND	Ground	**B15**	PETN0	Input	DPU SoC RX
A16	PERP0	Output	DPU SoC TX	**B16**	GND	GND	Ground
A17	PERN0	Output	DPU SoC TX	**B17**	RSVD		Ground
A18	GND	GND	Ground	**B18**	GND	GND	Ground
A19	RSVD		Not connected	**B19**	PETP1	Input	DPU SoC RX
A20	GND	GND	Ground	**B20**	PETN1	Input	DPU SoC RX
A21	PERP1	Output	DPU SoC TX	**B21**	GND	GND	Ground
A22	PERN1	Output	DPU SoC TX	**B22**	GND	GND	Ground
A23	GND	GND	Ground	**B23**	PETP2	Input	DPU SoC RX
A24	GND	GND	Ground	**B24**	PETN2	Input	DPU SoC RX
A25	PERP2	Output	DPU SoC TX	**B25**	GND	GND	Ground
A26	PERN2	Output	DPU SoC TX	**B26**	GND	GND	Ground
A27	GND	GND	Ground	**B27**	PETP3	Input	DPU SoC RX
A28	GND	GND	Ground	**B28**	PETN3	Input	DPU SoC RX

A29	PERP3	Output	DPU SoC TX	B29	GND	GND	Ground
A30	PERN3	Output	DPU SoC TX	B30	RSVD		Not connected
A31	GND	GND	Ground	B31	RSVD		Not connected
A32	RSVD		Not connected	B32	GND	GND	Ground
A33	RSVD		Not connected	B33	PETP4	Input	DPU SoC RX
A34	GND	GND	Ground	B34	PETN4	Input	DPU SoC RX
A35	PERP4	Output	DPU SoC TX	B35	GND	GND	Ground
A36	PERN4	Output	DPU SoC TX	B36	GND	GND	Ground
A37	GND	GND	Ground	B37	PETP5	Input	DPU SoC RX
A38	GND	GND	Ground	B38	PETN5	Input	DPU SoC RX
A39	PERP5	Output	DPU SoC TX	B39	GND	GND	Ground
A40	PERN5	Output	DPU SoC TX	B40	GND	GND	Ground
A41	GND	GND	Ground	B41	PETP6	Input	DPU SoC RX
A42	GND	GND	Ground	B42	PETN6	Input	DPU SoC RX
A43	PERP6	Output	DPU SoC TX	B43	GND	GND	Ground
A44	PERN6	Output	DPU SoC TX	B44	GND	GND	Ground
A45	GND	GND	Ground	B45	PETP7	Input	DPU SoC RX

A46	GND	GND	Ground	B46	PETN7	Input	DPU SoC RX
A47	PERP7	Output	DPU SoC TX	B47	GND	GND	Ground
A48	PERN7	Output	DPU SoC TX	B48	RSVD		Not connected
A49	GND	GND	Ground	B49	GND	GND	Ground
A50	RSVD		Not connected	B50	PETP8	Input	DPU SoC RX
A51	GND	GND	Ground	B51	PETN8	Input	DPU SoC RX
A52	PERP8	Output	DPU SoC TX	B52	GND	GND	Ground
A53	PERN8	Output	DPU SoC TX	B53	GND	GND	Ground
A54	GND	GND	Ground	B54	PETP9	Input	DPU SoC RX
A55	GND	GND	Ground	B55	PETN9	Input	DPU SoC RX
A56	PERP9	Output	DPU SoC TX	B56	GND	GND	Ground
A57	PERN9	Output	DPU SoC TX	B57	GND	GND	Ground
A58	GND	GND	Ground	B58	PETP10	Input	DPU SoC RX
A59	GND	GND	Ground	B59	PETN10	Input	DPU SoC RX
A60	PERP10	Output	DPU SoC TX	B60	GND	GND	Ground
A61	PERN10	Output	DPU SoC TX	B61	GND	GND	Ground
A62	GND	GND	Ground	B62	PETP11	Input	DPU SoC RX

A63	GND	GND	Ground	**B63**	PETN11	Input	DPU SoC RX
A64	PERP11	Output	DPU SoC TX	**B64**	GND	GND	Ground
A65	PERN11	Output	DPU SoC TX	**B65**	GND	GND	Ground
A66	GND	GND	Ground	**B66**	PETP12	Input	DPU SoC RX
A67	GND	GND	Ground	**B67**	PETN12	Input	DPU SoC RX
A68	PERP12	Output	DPU SoC TX	**B68**	GND	GND	Ground
A69	PERN12	Output	DPU SoC TX	**B69**	GND	GND	Ground
A70	GND	GND	Ground	**B70**	PETP13	Input	DPU SoC RX
A71	GND	GND	Ground	**B71**	PETN13	Input	DPU SoC RX
A72	PERP13	Output	DPU SoC TX	**B72**	GND	GND	Ground
A73	PERN13	Output	DPU SoC TX	**B73**	GND	GND	Ground
A74	GND	GND	Ground	**B74**	PETP14	Input	DPU SoC RX
A75	GND	GND	Ground	**B75**	PETN14	Input	DPU SoC RX
A76	PERP14	Output	DPU SoC TX	**B76**	GND	GND	Ground
A77	PERN14	Output	DPU SoC TX	**B77**	GND	GND	Ground
A78	GND	GND	Ground	**B78**	PETP15	Input	DPU SoC RX
A79	GND	GND	Ground	**B79**	PETN15	Input	DPU SoC RX

A80	PERP15	Output	DPU SoC TX	**B80**	GND	GND	Ground
A81	PERN15	Output	DPU SoC TX	**B81**	PRSNT2#	Output	x16 PCIe card detection, connected to PRESNT1#
A82	GND	GND	Ground	**B82**	RSVD		Not connected

Power Sequencing

The BF2500 Controller Card requires the 12V ATX PCIe power to come up prior or at the same time to the +12V at the PCIe edge connector.

NC-SI Management Interface

Pin#	SignalName	I/O	Description	Comment
1	REF_CLK	Input	50M REF CL	

Pin#	SignalName	I/O	Description	Comment
			K for NCSI BUS	
2	GND	GND	Ground	
3	ARB_IN	Input	NCSI hardware arbitration input	
4	GND	GND	Ground	

Pin#	SignalName	I/O	Description	Comment
5	ARB_OUT	Output	NCSI hardware arbitration output	
6	GND	GND	Ground	
7	RX_D0	Output	Receive data	Output for SoC
8	GND	GND	Ground	

Pin#	SignalName	I/O	Description	Comment
9	RX_D1	Output	Receive data	Output for SoC
10	GND	GND	Ground	
11	CRS_DV	Output	Carrier sense/ Receive Data Valid	
12	GND	GND	Ground	
13	TX_D0	Input	Trans	Input

Pin#	SignalName	I/O	Description	Comment
			mit data	for SoC
14	GND	GND	Ground	
15	TX_D1	Input	Transmit data	Input for SoC
16	GND	GND	Ground	
17	TX_EN	Input	Transmit enable	

Pin#	SignalName	I/O	Description	Comment
18	GND	GND	Ground	
19	NC			
20	NC			
21	I2C_SDA	Bidirectional	I2C Serial Data	GW_ARM1
22	GND	GND	Ground	
23	I2C_SCL	Bidirectional	I2C Serial Clock	GW_ARM1
24	GND	GND	Ground	
25	GND	GND	Ground	

Pin#	SignalName	I/O	Description	Comment
26	GND	GND	Ground	
27	NC			
28	UART_TX	Output	Transmit data	Output for SoC
29	NC			
30	UART_RX	Input	Receive data	

External Power Supply Connector (6-pin ATX 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 6-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 6-pin ATX type connector can have different pinouts.

Pin#	Signal Name	Description
1	12V	ATX Supplied 12V
2	12V	ATX Supplied 12V
3	12V	ATX Supplied 12V
4	GND	Power Return
5	GND	Power Return
6	GND	Power Return

External +12V Mechanical Pinout Diagram

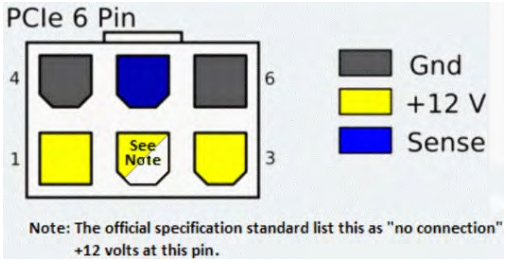

Hardware Installation

Installation and initialization of the BlueField-2 BF2500 DPU Controller require attention to the mechanical, power, and precautions for rack-mounted equipment.

Safety Warnings

> Safety warnings are provided here in the English language. For safety warnings in other languages, refer to the DPU Controller Installation Safety Instructions document available on the public website.

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

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

Installation Procedure Overview

The installation procedure of BlueField-2 DPU Controller involves the following steps:

Step	Procedure	Direct Link
1	Unpack the package and confirm that you have received all the required components	Refer to Package Contents
2	Check the system's hardware and software requirements.	Refer to System Requirements
3	Pay attention to the airflow consideration within the JBOF system	Refer to Airflow Requirements
4	Follow the pre-installation checklist	Refer to Pre-Installation Checklist
5	Install the BF2500 DPU Controller in the JBOF system	Refer to Installation Instructions

7	Connect the external power supply to the BF2500 DPU Controller
8	Power on the system
9	Perform initial system bring-up

System Requirements

Hardware Requirements

> Unless otherwise specified, NVIDIA 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.

A system that supports PCI Express high power cards is required for installing the card.

The system should be able to support 75W through the PCIe x16 interface and an additional 75W through the ATX PCIe power connector

Airflow Requirements

BlueField-2 BF2500 DPU Controller is offered with one airflow pattern: from the BlueField-2 SoC to the network ports.

> All systems in the same rack should be planned with the same airflow direction. All components need to have the same airflow direction.

Please refer to the Specifications section for airflow numbers for each specific card model.

Software Requirements

- See Operating System under System Requirements section under the Introduction section.
- Software Stacks - The BF2500 DPU Controller is shipped with Linux based Operating System burned on it which includes all needed drivers. For more information, please refer to the BlueField-2 Software User Manual.
- Check System Requirements Overview under Introduction for more details.

Safety Precautions

The BF2500 card 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.

Pre-Installation Checklist

1. Unpack the BF2500 DPU Controller

Unpack and remove the BF2500 DPU Controller. 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 cards must be placed on an antistatic surface. For package contents please refer to Package Contents.

> Please note that if the card is removed hastily from the antistatic bag, the plastic ziplock may harm the EMI fingers on the QSFP56 connector. Carefully remove the card from the antistatic bag to avoid damaging the EMI fingers.

2. Turn off the power to the JBOF system.

Turn off the power to the JBOF system, and disconnect the power cord and remove the cover. Refer to the JBOF system documentation for instructions. Before you install the BF2500 DPU Controller, make sure that the system is disconnected from power and any networks.

Installation Instructions

This section provides detailed instructions on how to install your BlueField-2 2500 DPU Controller in a system.

> The BlueField-2 BF2500 DPU Controller should be installed only in a JBOF System as it functions as a PCIe root-complex (RC) initiating PCIe bus operations. Installing it in a regular host system may damage the card.

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

Step 1. Open the system case.

Step 2. Locate an available PCI Express x16 slot.

Step 3. Applying even pressure at both corners of the card, insert the BF2500 DPU Controller in a PCI Express slot until firmly seated.

Step 4. Secure the bracket to the system with the bracket screw.

Step 5. Close the system case

Cables and Modules

Networking Cables

1. All networking 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 adapter card. Do this by using a cable holder or tying the cable to the rack.
    b. Determine the correct orientation of the connector to the card before inserting the connector. Do not try and insert the connector upside down. This may damage the adapter card.

c. Insert the connector into the adapter card. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the adapter card.
d. Verify 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 card.

3. After inserting a cable into a port, the Amber 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).
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. See Networking LED Interfaces.
5. Care should be taken as 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. The LED indicator will turn off when the cable is unseated.

RJ45 Cable for the 1GbE OOB Management Interface

NOTE: A CAT5 cable should be used.

USB Interface Cabling

Debugging and loading new versions of the operating systems and firmware requires the use of Type A to Type B mini-USB 2.0 Cable. The following figure shows an example of a cable with a USB Type A connector on one end an a mini-USB Type B connector on the other.

Bring-Up and Driver Installation

⚠ It is recommended to upgrade your BlueField product to the latest software and firmware versions in order to enjoy the latest features and bug fixes.

BlueField Software

NVIDIA provides software that enables users to fully utilize the BlueField® DPU and enjoy the rich feature-set it provides. Using BlueField software packages, users are able to:

- Quickly and easily boot an initial Linux image on your development board
- Port existing applications to and develop new applications for BlueField
- Patch, configure, rebuild, update or otherwise customize your image
    - Debug, profile, and tune their development system using open source development tools taking advantage of the diverse and vibrant Arm ecosystem.

The BlueField family of DPU devices combines an array of 64-bit Arm v8 A72 cores coupled with the ConnectX® interconnect. Standard Linux distributions run on the Arm cores allowing common open source development tools to be used. Developers should find the programming environment familiar and intuitive which in turn allows them to quickly and efficiently design, implement and verify their control-plane and data-plane applications.

BlueField SW ships with the NVIDIA BlueField Controller Cards. BlueField SW is a reference Linux distribution based on the Yocto Poky distribution and extended to include the NVIDIA OFED stack for Arm and a Linux kernel which supports NVMe-oF. This SW distribution is capable of running all customer-based Linux applications seamlessly. Yocto also provides an SDK that contains an extremely flexible cross-build environment allowing software targeted for the BlueField DPU to build on virtually any x86 server running any Linux distribution.

The following are other software elements delivered with BlueField DPU:

- Arm Trusted Firmware (ATF) for BlueField
- UEFI for BlueField
- Hardware Diagnostics
- NVIDIA OFED stack
- NVIDIA MFT

For the BlueField software and software user manual, please contact your local support.

Before installing the preferred OS on the BlueField Controller Card, make sure you install the card only in a JBOF system. Installing it in a host system may damage the card.

Software On eMMC

The BlueField Controller Card boots off eMMC upon power-up. The image flashed on the eMMC from the factory is the Yocto Linux.

Run the following command to discover the BlueField Software version:

```
cat /etc/bluefield_version
```

Yocto Distribution Installation

The BlueField tarball comes with pre-built Yocto images that can be installed. Please refer to The BlueField Software User Manual for the basic Yocto Installation.The core-image-full image is a full root filesystem image that is appropriate for imaging on the rootfs partition of the eMMC. Refer to the BlueField Software User Manual for instructions on how to build it yourself, or contact your local support to get the image.

1. Prepare the host environment. For more information, refer to Preparing the Host-Side Environment.
2. Boot the BlueField Controller Card over USB using the samples/install.bfb image.
3. Refer to samples/README.install for instructions.
4. To boot the BlueField Controller Card over USB from the server host, run:

   ```
   cat install.bfb > /dev/rshim0/boot
   ```

5. To prepare the eMMC for Yocto installation from the Arm, run:

   ```
   /opt/mlnx/scripts/bfinst --fullfs /tmp/core-image-full-bluefield.tar.xz
   ```

6. After the installation is done, execute lowercase reboot on the Arm.

```
shutdown -r now
```

7. Verify the version via:

```
cat /etc/bluefield_version
```

PXE Server Configuration on Host Side

Before installing CentOS 7 on the BlueField Controller Card, you need to configure the PXE server on the host side (x86) to allow the deployment of the CentOS image over the BlueField Controller Card.

Download the CentOS installation iso file from the following link:

```
# Download the centos installation iso file from http://mirror.centos.org/altarch/7/isos/aarch64/CentOS-7-aarch64-Everything.iso
# cd <BF_INST_DIR>/distro/rhel/pxeboot
# ./setup.sh -d <BF_INST_DIR> -i <centos-installation.iso> [-c <ttyAMA0 >]
```

⚠ UART0 (ttyAMA0) is used by default, or you can use "-c ttyAMA0" to manually specify UART0.

Installing Linux on BlueField Controller Card

This section demonstrates CentOS 7.4 installation on the BlueField Controller Card. Other OSs work similarly with the PXE boot installation process.

Before installing the preferred OS on the BlueField Controller Card, make sure you install the BlueField Controller Card in a JBOF System. Installing it in a host system may damage the card.

Software Requirements

- CentOS 7.4 Linux OS. To get CentOS 7.4 image, run:

```
wget http://archive.kernel.org/centos-vault/altarch/7.4.1708/isos/aarch64/CentOS-7- aarch64-Everything.iso
```

> ⚠ Some required drivers do not compile and load if running CentOS 5.x or earlier.

> ⚠ Please note that CentOS 7.5 is not supported.

- Access to the latest BlueField Controller Card SW bundle:
  NVIDIA uses box.com to distribute BlueField software. Contact your sales/support representative for a custom link to download BlueField software releases.
- In this document, we assume the tarball BlueField-1.0.alphaX.XXXXX.tar.gz is extracted at / root, to do this, run the following command:

```
tar -xvf BlueField-1.0.alphaX.XXXXX.tar.xz -C /root
```

Preparing Host-Side Environment

Before installing the preferred OS on the BlueField Controller Card, the host must be set up for it to be capable of provisioning the BlueField Controller Card. The RShim USB driver is installed on the host to communicate with the RShim device on the BlueField DPU. The RShim USB driver must be installed so that it can push the initial bootloader and supply the OS image for PXE boot through the USB connection.

> ⚠ This process only needs to be done on the host machine which is provisioning the BlueField Controller Card, it is not required on the end machine.

Setup Procedure With Installation Script

If the host is running CentOS 7 (or equivalent) on the host, you may run a script to complete all the steps detailed in Preparing the Host-Side Environment.

```
/root/BlueField-1.0.alphaX.XXXXX/distro/rhel/pxeboot/setup.sh    \
-d /root/BlueField-1.0.alphaX.XXXXX/ \
-i /root/CentOS-7-aarch64-Everything.iso \
-o /root/dd-rhel7.4-mlnx-ofed-4.2-1.4.10.0-aarch64.iso \
-c ttyAMA0\
-k
```

Note that there should be no firewall blocking the IP communication between the BlueField Controller Card and the server host machine. If a firewall exists, disable it with the following commands:

```
iptables -F
iptables -t
nat -F
```

- The “-d” flag points to where the tar file has been extracted from, the script uses this directory to find all the source code it needs.
- The “-i” flag points to the OS installation disk. This is the image that is accessed via PXE boot to install the OS on the BlueField Controller Card.
- The “-o” flag points to the NVIDIA OFED driver disk for Arm. Download and extract it from http://www.mellanox.com/page/products_dyn?product_family=34.
- The “-c” flag specifies the default UART port for the OS to use since the BlueField DPU has two Arm UARTs. For the BlueField Controller Card, “ttyAMA0” is used, which is UART0.
- The “-t” flag is optional and needed for nonpxe boot. When specified and given the argument of what Controller card is set (BlueField Controller Card in this case), it generates a “nonpxe.bfb” file which contains the install kernel and rootfs. If this file is pushed to the RShim boot device, it automatically runs the installation process and skips the initial UEFI PXE boot operations. (the -t flag). Please refer to distro/rhel/pxe/README.
- The optional “-k” flag kickstarts auto-installation based on a default kickstart file which is installed as /var/pxe/ks/ks.cfg (optional).

Setup Procedure Without Installation Script

If the host is running CentOS 7 or equivalent, please refer to Preparing the Host-Side Environment for a simpler way to perform the installation using an installation script.

The following sections demonstrate CentOS 7 installation, however, installation in other environments should be relatively similar.

Step 1: Set up RShim Interface

The RShim driver communicates with the RShim device on the BlueField DPU. The RShim is in charge of many miscellaneous functions of the DPU, including resetting the Arm cores, providing the initial bootstream, and using the TMFIFO and the RShim network, to exchange network and console data with the host.

The RShim device can be reached by the host via the USB connector.

Step 2: Install RShim Drivers

To install the kernel modules, please follow the instruction in section RShim Host Driver.

To install the kernel modules, please follow the instruction in NVIDIA BlueField DPU Family Software Documentation under BlueField Software Overview > Installation and Initialization > RShim Host Driver.

Step 3: Configure TFTP Server

The host should be configured to act as a TFTP server to the BlueField Controller Card via the USB RShim network. This server provides the required files by the BlueField Controller Card to perform the PXE boot for installing the preferred OS.

> ⚠ Configuring the TFTP server requires a TFTP package. If it is not installed, install it via “yum install tftp” or “apt-get tftp”, depending on your Linux distribution.
> Note: On some versions, the TFTP package cannot be found. In such cases, install “xinetd”.

1. Extract the OS image and copy the required PXE boot components:

```
mount -t iso9660 -o loop CentOS-7-aarch64-Everything.iso /mnt
mkdir -p /var/lib/tftpboot/centos/7.4
cp /mnt/EFI/BOOT/BOOTAA64.EFI /var/lib/tftpboot/ cp /mnt/EFI/BOOT/grubaa64.efi /var/lib/tftpboot/
cp /mnt/images/pxeboot/vmlinuz /var/lib/tftpboot/centos/7.4
cp /mnt/images/pxeboot/initrd.img /var/lib/tftpboot/centos/7.4/initrd-orig.img
```

2. Patch the initrd with the eMMC driver and TMFIFO (RShim network) driver:

```
mkdir -p /tmp/.bfcentos mkdir -p $/tmp/.bfinstdd cd /tmp/.bfcentos
xzcat /var/lib/tftpboot/centos/7.4/initrd-orig.img | cpio -idm mount
/root/BlueField-1.0.alpha3.10409/distro/rhel/bluefield_dd/bluefield_dd-4.11.0- 22.el7a.aarch64.iso
/tmp/.bfinstddmkdir -p usr/lib/modules/4.11.0-22.el7a.aarch64/updates/cp
/tmp/.bfinstdd/lib/modules/4.11.0-22.el7a.aarch64/updates/dw_mmc*.ko usr/lib/mod- ules/
4.11.0-22.el7a.aarch64/updates/cp
/tmp/.bfinstdd/lib/modules/4.11.0-22.el7a.aarch64/updates/tmfifo.ko usr/lib/mod- ules/4.11.0-22.el7a.aarch64/
updates/cp
/root/BlueField-1.0.alpha3.10409/distro/rhel/bluefield_dd/bluefield_dd-4.11.0- 22.el7a.aarch64.iso ./
bluefield_dd.iso
umount /tmp/.bfinstdd; rmdir /tmp/.bfinstdd chown root:root
* -R
depmod -b
/tmp/.bfcentos 4.11.0-22.el7a.aarch64
find . | cpio
-oc | xz --check=crc32 --lzma2=dict=32MiB >
/var/lib/tftpboot/centos/7.4/ initrd.img
```

> ⚠ These commands assume that you are using kernel version “4.11.0-22.el7a.aarch64”. If you are using a different version, utilize the corresponding bluefield_dd.iso. If none is found, compile one by running the following: source /path/to/SDK/environment-setup-aarch64-poky-linux; unset LDFLAGS; ./build- dd.sh /path/to/kernel-devel-4.11.0-44.el7a.aarch64.rpm

3. Change the grub configuration to PXE boot over the right location:

```
cat >/var/lib/tftpboot/grub.cfg <<EOF
menuentry 'Install centos/7.4 AArch64 - BlueField'
--class red --class
gnu-linux
--class gnu --class os {
linux (tftp)/centos/7.4/vmlinuz ro ip=dhcp method=http://192.168.100.1/centos7 inst.dd=/bluefield_dd.iso
console=ttyAMA0
initrd (tftp)/centos/7.4/initrd.img
} EOF
```

4. Start the TFTP server:

```
systemctl restart tftp
```

> ⚠ Based on the system, the user may need to use “system TFTP restart” instead. Also, if required, the user might need to switch use “xinetd” instead of “TFTP”.

Step 4: Set Up the DHCP Server

DHCP server set up on the host is required for BlueField Controller Card to get a private IP from the host for PXE boot process completion. Configure the correct server names and domain names so that the BlueField Controller Card can connect to the network via the host later on.

1. Get the server/domain names on the host:

   ```
   bash-4.2$ cat /etc/resolv.conf
   # Generated by NetworkManager search internal.mlnx.com labs.mlnx
   nameserver 10.15.2.29
   nameserver 10.15.2.16
   ```

   This example shows that the domains are internal.mlnx.com and labs.mlnx, and the names of the servers are 10.15.2.29 and 10.15.2.16.
2. Set up the DHCP config file accordingly:

   ```
   cat >/etc/dhcp/dhcpd.conf <<EOF
   allow booting;
   allow bootp;
   subnet 192.168.100.0 netmask 255.255.255.0 {
   range 192.168.100.10 192.168.100.20;
   option broadcast-address 192.168.100.255;
   option routers 192.168.100.1;
   option domain-name-servers 10.15.2.29 10.15.2.16;
   # Set the domain search according to the network configuration option domain-search "internal.tilera.com"
   "mtbu.labs.mlnx"; next-server 192.168.100.1;
   filename "/BOOTAA64.EFI";
   }
   ```

```
# Specify
the IP address for
this client. host pxe_client {
hardware ethernet 00:1a:ca:ff:ff:01; fixed-address 192.168.100.2;
} EOF
```

⚠ It is recommended to back up the previous dhcpd.conf file before overwriting it.

Step 5: Set Up the HTTP Server

The TFTP server allows the PXE boot to load the initrd and kernel. The BlueField Controller Card obtains all the other required sources through the network, thus, making it necessary to set up an HTTP.

⚠ Setting up the HTTP server requires the HTTP package. If it is not installed, please install it via “yum install httpd” or “apt-get httpd”, depending on your Linux distribution.

To configure the http server to serve the contents of the installation disk, run the following command:

```
cat >/etc/httpd/conf.d/pxeboot.conf <<EOF Alias /centos7
/mnt
<Directory /mnt>
    Options Indexes FollowSymLinks Require ip 127.0.0.1
    192.168.100.0/24
</Directory>
EOF
systemctl enable httpd systemctl restart
httpd
```

Flashing BlueField Controller Card Bootloader Code

Before installing an OS, flash the bootloader code first. The BlueField Controller Card is shipped with an initial bootloader code, and should be updated with the following instructions.

Opening Terminal Connection to BlueField Controller Card

To open a console window to the BlueField Controller Card, a terminal application is required. The application “minicom” is used for the flow, however, any standard terminal application can work, e.g. “screen”.

> ⚠ Install minicom by running “yum install minicom” or “apt-get install minicom”.

1. On the host, type “minicom” to open minicom on the current terminal, use “minicom -s” to set it up.
2. Go to the settings menu by pressing “Ctrl-a + o” (the setting menu opens by default when launching with the “-s” option). Navigate to the “Serial port setup” submenu and set the “Serial Device” to the one connected (should be one of the /dev/ttyUSBx if using the serial- UART cable).
3. Change the baud rate to 115200 8N1, and ensure that the hardware and software flow control are set to “No”.
   Minicom Settings - Example

4. Select “Save setup as dfl” in order not to have to set it again in the future.

Using Initial Install Bootstream

1. On the host side, ensure that the RShim driver is running:

   ```
   $ systemctl status rshim
   ```

   An RShim device is located under the /dev directory, if you only have one, it should be “rshim0”:

```
[root@bu-lab02 ~]# ls /dev/rshim0/
boot      console net       rshim
```

The boot device is used to push the bootstream to the BlueField Controller Card. Upon writing to it, it automatically resets the Arm cores so that it is booted using the pushed bootstream.

> ⚠ The console device can be used as a console instead of the serial-USB console. The primary bootloader does not support this device, however, UEFI and Linux support it. In cases where the special UART adapter board is unavailable, this can be used instead.

2. Push the initial install bootstream to the BlueField Controller Card:

```
cat /root/BlueField-1.0.alphaX.XXXXX/sample/install.bfb > \
/dev/rshim0/boot
```

On the terminal, various boot messages appear until Linux is loaded. This is the Yocto embedded Linux running off the kernel initramfs pushed in the bootstream.

3. When prompted, type in “root” to get to the command prompt without any password.
Yocto Log

```
done.
Starting OpenBSD Secure Shell server: sshd
done.
Starting rpcbind daemon...done.
starting statd: done
exportfs: can't open /etc/exports for reading
NFS daemon support not enabled in kernel
Starting syslogd/klogd: done

Poky (Yocto Project Reference Distro) 2.3.1 bluefield /dev/ttyAMA0

bluefield login: root
root@bluefield:~#
CTRL-A Z for help |115200 8N1 | NOR | Minicom 2.6.2  | VT102 |      Offline
```

4. After Linux is loaded, in the terminal, run the /opt/mlnx/scripts/bfrec script to update the bootloader.

Installing CentOS 7.4 on BlueField Controller Card

⚠ If the error “no root is found” appears in the installation process, check or disable the firewall as needed on the server host machine.

Full PXE Boot Installation

1. Get to the UEFI boot menu.
   a. Reboot the BlueField Controller Card by typing “reboot” on the console. A “UEFI firmware…” message should appear and the screen clears.
   b. Press ESC several times until you enter the UEFI boot menu.

      UEFI Boot Menu

   c. On the host, restart the DHCP and TFTP service:

```
systemctl restart dhcpd
systemctl restart tftp #might be xinetd
```

d. Navigate to the Boot Manager.
UEFI Boot Manager

e. Select EFI Network, it will then use the TFTP service on the host to discover all available PXE boot options. Shortly after, a "..Fetching Netboot Image" message will appear enabling CentOS installation.
Option to Install CentOS

f. Select CentOS download.

> ⚠ This process may take few minutes as it fetches data over the USB network. Running "ifconfig" on the host and monitoring the RX/TX packets on the "tmfifo_net0" network indicates that the fetching data process is not complete.

g. Follow the installation instructions in the configuration menu. Recommended settings are included.

⚠ These configuration inputs are not needed when the kickstart option “-k” is specified when running the setup.sh script.

```
================================================================================
================================================================================ VNC

Text mode provides a limited set of installation options. It does not offer custom partitioning for full
control over the disk layout. Would you like to use VNC mode instead?

1)  Start VNC

2)  Use text mode

Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: 2
================================================================================
================================================================================
Installation:main* 2:shell 3:log 4:storage-lo> Switch tab: Alt+Tab | Help: F1
1) [x] Language settings    2) [!] Time settings
(English (United States))   (Timezone is not set.)
3) [!] Installation source  4) [!] Software selection (Processing...)      (Processing...)
5) [!] Installation Destination 6) [x] Kdump
(No disks selected) (Kdump is enabled)
7) [x] Network configuration    8) [!] Root password
(Wired (eth0) connected)    (Password is not set.)
9) [!] User creation
(No user will be created)
Please make your choice from above ['q' to quit | 'b' to begin installation | 'r' to refresh]: 2
================================================================================
================================================================================
Time settings Timezone: not set
NTP servers:not configured

1)  Set timezone
2)  Configure NTP servers
Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: 1
================================================================================
================================================================================
Timezone settings Available regions
Available regions
```

```
1)  Europe  6)  Pacific     10) Arctic
2)  Asia    7)  Australia   11) US
3)  America 8)  Atlantic    12) Etc
4)  Africa  9)  Indian
5)  Antarctica

Please select the timezone.
Use numbers or type names directly [b to region list, q to quit]: 11
================================================================================
================================================================================
Timezone settings

Available timezones in region US
1)  Alaska  4) Eastern  6) Mountain
2)  Arizona 5) Hawaii   7) Pacific
3)  Central
Please select the timezone.
Use numbers or type names directly [b to region list, q to quit]: 4
================================================================================
================================================================================
Installation

1) [x] Language settings    2) [x] Time settings (English (United States))      (US/Eastern timezone)
3) [x] Installation source  4) [x] Software selection (http://192.168.100.1/centos7)        (Minimal
Install)
5) [!] Installation Destination 6) [x] Kdump
(No disks selected) (Kdump is enabled)
7) [x] Network configuration    8) [!] Root password
(Wired (eth0) connected)    (Password is not set.)
9) [!] User creation
(No user will be created)
Please make your choice from above ['q' to quit | 'b' to begin installation | 'r' to refresh]: 4
================================================================================
================================================================================
Base environment Software selection

Base environment

1)  [x] Minimal Install 6) [ ] Server with GUI
2)  [ ] Compute Node    7) [ ] GNOME Desktop
3)  [ ] Infrastructure Server   8) [ ] KDE Plasma Workspaces
```

```
4)  [ ] File and Print Server   9) [ ] Development and Creative
5)  [ ] Basic Web Server    Workstation
Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: 9
================================================================================
================================================================================
Base environment Software selection

1)  [ ] Minimal Install 6)  [ ] Server with GUI
2)  [ ] Compute Node    7)  [ ] GNOME Desktop
3)  [ ] Infrastructure Server   8)  [ ] KDE Plasma Workspaces
4)  [ ] File and Print Server   9)  [x] Development and Creative
5) [ ] Basic Web Server Workstation
Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: c
================================================================================
================================================================================
Installation

1) [x] Language settings    2) [x] Time settings (English (United States))      (US/Eastern timezone)
3) [!] Installation source  4) [!] Software selection (Processing...)       (Processing...)
5) [!] Installation Destination 6) [x] Kdump
(No disks selected) (Kdump is enabled)
7) [x] Network configuration    8) [!] Root password
(Wired (eth0) connected)    (Password is not set.)
9) [!] User creation
(No user will be created)
Please make your choice from above ['q' to quit | 'b' to begin installation | 'r' to refresh]: 5
================================================================================
================================================================================
Probing storage...
Installation Destination

[x] 1) : 13.75 GiB (mmcblk0)

1 disk selected; 13.75 GiB capacity; 1007.5 KiB free ...

Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: c
================================================================================
================================================================================
Autopartitioning Options

[ ] 1) Replace Existing Linux system(s) [x] 2) Use All Space
```

```
[ ] 3) Use Free Space

Installation requires partitioning of your hard drive. Select what space to use for the install target.

Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: c
================================================================================
================================================================================
Partition Scheme Options [ ] 1) Standard Partition [ ] 2) Btrfs
[x] 3) LVM
[ ] 4) LVM Thin Provisioning

Select a partition scheme configuration.

Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: 1
================================================================================
================================================================================
Partition Scheme Options [x] 1) Standard Partition [ ] 2) Btrfs
[ ] 3) LVM

[ ] 4) LVM Thin Provisioning

Select a partition scheme configuration.

Please make your choice from above ['q' to quit | 'c' to continue | 'r' to refresh]: c Generating
updated storage configuration
Checking storage configuration...

================================================================================
================================================================================
Installation

1)  [x] Language settings   2) [x] Time settings (English (United States))      (US/Eastern timezone)
3) [x] Installation source  4) [x] Software selection (http://192.168.100.1/centos7)        (Development
and Creative
5) [x] Installation Destination     Workstation) (Automatic partitioning    6) [x] Kdump
selected)   (Kdump is enabled)
7) [x] Network configuration    8) [!] Root password
(Wired (eth0) connected)    (Password is not set.)
9) [!] User creation
(No user will be created)
Please make your choice from above ['q' to quit | 'b' to begin installation | 'r' to refresh]: 8
```

```
================================================================================
================================================================================
Please select new root password. You will have to type it twice.

Password:
Password (confirm):
================================================================================
================================================================================
Question

The password you have provided is weak: The password fails the dictionary check
-   it is based on a dictionary word. Would you like to use it anyway?

Please respond 'yes' or 'no': yes
================================================================================
================================================================================
Installation

1)  [x] Language settings   2) [x] Time settings (English (United States))      (US/Eastern timezone)
3) [x] Installation source  4) [x] Software selection (http://192.168.100.1/centos7)        (Development
and Creative
5) [x] Installation Destination     Workstation) (Automatic partitioning    6) [x] Kdump
selected)   (Kdump is enabled)
7) [x] Network configuration    8) [x] Root password (Wired (eth0) connected)       (Password is set.)
9) [ ] User creation
(No user will be created)
Please make your choice from above ['q' to quit | 'b' to begin installation | 'r' to refresh]: b
```

h. Enter “b” and press “enter” to initiate the installation process.
i. Press “Enter” to reboot into CentOS.

CentOS Installation Completion Screen

Non-PXE Boot Installation

When the setup script is run with the "-t" option, it generates a nonpxe.bfb file at the directory where the script is run. The directory contains the install kernel and rootfs which are usually loaded by UEFI during the initial PXE boot stage. Thus, if pushing this file, the host TFTP server no longer needs to be used and UEFI would automatically load the install kernel and rootfs from the boot FIFO. Together with the "-k" kickstart option, the host can be configured to initiate non-PXE boot and automatic CentOS installation, as long as the host HTTP and DHCP servers are working. To kick off the installation process, run the following command on the host:

```
cat nonpxe.bfb > /dev/rshim0/boot; sleep 2; systemctl restart dhcpd
```

MLNX_OFED Installation

⚠ This section is relevant to non-Yocto Operating Systems only.

Installing MLNX_OFED on Arm Cores

Prerequisite Packages for Installing MLNX_OFED

- MLNX_OFED installation requires some prerequisite packages to be installed on the system.
  Currently, CentOS installed on the BlueField Controller Card has a private network to the host via the USB connection, and it can be used to Secure Copy Protocol (SCP) all the required packages. However, it is recommended for the BlueField Controller Card to have a direct access to the network to use “yum install” to install all the required packages. For direct access to the network, set up the routing on the host via:

```
iptables -t nat -o eth0 -A POSTROUTING -j MASQUERADE
echo 1 > /proc/sys/net/ipv4/ip_forward
systemctl restart dhcpd
```

  ⚠ “eth0” is the outgoing network interface on the host. Change this according to your system requirements.

  ⚠ These commands are not saved in Linux startup script, and might be needed to be applied again after host machine reboots.

- Reset the BlueField Controller Card network for Internet connection (access to the web) as long as the host is connected:

```
[root@localhost
~]# ifdown eth0;
ifup eth0 [root@localhost
~]# ping google.com
PING google.com (172.217.10.142) 56(84)
bytes of data.
64 bytes from lga34s16-in-f14.1e100.net (172.217.10.142): icmp_seq=1 ttl=53 time=19.2 ms
64 bytes from lga34s16-in-f14.1e100.net (172.217.10.142): icmp_seq=2 ttl=53 time=17.7 ms
64 bytes from
lga34s16-in-f14.1e100.net (172.217.10.142): icmp_seq=3 ttl=53 time=15.8 ms
```

- Run “yum install” to install all the required MLNX_OFED packages:

```
yum install rpm-build
yum group install "Development Tools" yum install kernel-devel-`uname -r`
yum install valgrind-devel libnl3-devel python-devel yum install tcl tk
```

Note that this is not needed if you installed CentOS 7 with the kickstart ("-k") option.
/auto/sw_mc_soc_project/distro/rhel/kernel-devel-4.11.0-22.el7a.aarch64.rpm

Removing Pre-installed Kernel Module

There are cases where the kernel is shipped with an earlier version of the mlx5_core driver taken from the upstream Linux code. This version does not support the BlueField Arm, but is loaded before the MLNX_OFED driver, and therefore, needs to be removed.

To remove the kernel module from the initramfs, run the following command:

```
mkdir /boot/tmp
cd /boot/tmp
gunzip < ../initramfs-4*64.img | cpio -i
rm -f lib/modules/4*/updates/mlx5_core.ko
rm -f lib/modules/4*/updates/tmfifo*.ko
cp ../initramfs-4*64.img ../initramfs-4.11.0-22.el7a.aarch64.img-bak
find | cpio -H newc -o | gzip -9 > ../initramfs-4*64.img
rpm -e mlx5_core
depmod -a
```

Installing MLNX_OFED on the BlueField Controller Card

1. Copy the MLNX_OFED image to the BlueField Controller Card via the USB network. The MLNX_OFED images should be provided in the software drop:

   ```
   scp MLNX_OFED_LINUX-4.2-1.4.8.0-rhel7.4alternate-aarch64.iso \
   root@192.168.100.2:/root
   ```

2. Mount the image on the BlueField Controller Card:

```
mount /root/MLNX_OFED_LINUX-4.2-1.4.8.0-rhel7.4alternate-aarch64.iso /mnt
```

3. Install MLNX_OFED.
   If the kernel on the BlueField is 4.11.0-22.el7a.aarch64, run:

```
cd /mnt
# ./mlnxofedinstall --bluefield
```

   If the kernel is different than 4.11.0-22.el7a.aarch64, run:

```
cd /mnt
# ./mlnxofedinstall --add-kernel-support --skip-repo
```

   ⚠ For OFED to support DPDK, use the arguments “--upstream-libs” and “--dpdk”.

```
./mlnxofedinstall --distro rhel7.4alternate --add-kernel-support --upstream-libs
--dpdk
```

   This step might take longer than expected to be completed. If you are using a different pack- age than the required one, run “yum install”.

   ⚠ If the date is not set correctly while installing MLNX_OFED, first, set the date (e.g date -s 'Mon Feb 5 15:02:10 EST 2018'), then run the installation.

4. Restart openibd:

```
/etc/init.d/openibd restart
```

Updating BlueField Controller Card Firmware

> ⚠ The below commands apply to MBF1M616A-CSNAT. The commands vary per OPN.

The below steps demonstrate how to manually update the firmware if the automatic process fails. The firmware image can be found in the BlueField Software package.

1. Copy the firmware image to the BlueField Arm:

```
scp fw-BlueField-rel--XX_XX_XXXX-MBF1M6X6A-CSNA_Ax.ini.bin \
root@192.168.100.2:/root
[root@localhost ~]# mst start
Starting MST (Mellanox Software Tools) driver set Loading MST PCI module - Success
Loading MST PCI configuration module - Success Create devices
Unloading MST PCI module (unused) - Success [root@localhost ~]# mst status
MST modules:
------------
MST PCI module is not loaded
MST PCI configuration module loaded
MST devices:
------------
/dev/mst/mt41682_pciconf0
- PCI configuration cycles access.
domain:bus:dev.fn=0000:04:00.0 addr.reg=88 data.reg=92
Chip revision is: 00
```

2. The output indicates that the device is “/dev/mst/mt41682_pciconf0”. To update the firmware:

```
flint -d /dev/mst/mt41682_pciconf0 b \
-i /root/ fw-BlueField-rel-XX_XX_XXXX-MBF1M6X6A-CSNA_Ax.ini.bin
```

When using the mlx and ini files, use the following command instead:

```
mlxburn -d /dev/mst/mt41682_pciconf0 -fw fw-BlueField.mlx -c bf.ini
```

To burn the firmware which comes with OFED after OFED is installed, run:

```
/opt/mellanox/mlnx-fw-updater/firmware/mlxfwmanager_sriov_dis    -force
```

3. Power cycle the BlueField Controller Card for the new firmware to take effect.

```
root@bluefield:~# flint -i /opt/fw-Bluefield-red-18_99_4608_MBF1M6X6A-CSNA_Ax-
Flexboot-3.5.404_UEFI-14.15.20.bin -d /dev/mst/mt41682_pciconf0 b

    Current FW version flash:     18.24.0013
    New FW version:               18.99.4608

Burning FW image without signatures - 55%
Burning FW image without signatures - OK
Restoring signature
-I- To load new FW run mlxfwreset or reboot machine.
root@bluefield:~#
```

> ⚠ After MLNX_OFED is installed on the Arm cores, use the mlx5_core driver to use the two Ethernet ports on the BlueField Controller Card. If the Ethernet ports on the BlueField Controller Card are connected to the network, there is no need to bridge the host via RShim net to access the network.

To install the kernel modules, please follow the instruction in section RShim Host Driver.

Troubleshooting

The BF2500 Controller Card no longer works	• Reseat the BF2500 Controller Card in its slot or a different slot, if necessary • Try using another cable • Reboot the system
BF2500 Controller Card stopped working after installing another BF2500 card	• Try removing and re-installing all BF2500 Controller Cards • Check that cables are connected properly
Link indicator light is off	• Try another port on the switch • Make sure the cable is securely attached • Check you are using the proper cables that do not exceed the recommended lengths • Verify that your switch and BF2500 Controller Card port are compatible
Link light is on, but with no communication established	• Check that both the BF2500 Controller Card and its link are set to the same speed and duplex settings

Specifications

Ensure your system can support 75W or greater system power supply through the PCIe x16 interface and an additional 75W through the supplementary 6-pin ATX power supply connector

MBF2H516B-CEEOT / MBF2H516B-CENOT Specifications

BlueField-2 DPU	BlueField-2 P-Series - 8 Cores - 550MHz/2750MHz	
Physical	**Card Dimensions (FHHL):** 2.71 in. x 6.6 in. (68.90mm x 167.65 mm) **Bracket Dimensions:** 4.73 in. x 0.72 in. (121.0mm x 18.4mm)	
	Connector: Dual QSFP56 (copper and optical)	
Protocol Support	**Ethernet**: 100GBASE-CR4, 100GBASE-KR4, 100GBASE-SR4, 50GBASE-R2, 50GBASE-R4, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-SR4, 40GBASE-LR4, 40GBASE-ER4, 40GBASE-R2, 25GBASE-R, 20GBASE-KR2, 10GBASE-LR,10GBASE-ER, 10GBASE-CX4, 10GBASE-CR, 10GBASE-KR, SGMII, 1000BASE-CX, 1000BASE-KX, 10GBASE-SR	
	Data Rate: 1/10/25/40/50/100 Gb/s Ethernet	
	PCI Express Gen 4.0: SERDES @16.0GT/s, 16 lanes (3.0, 2.0 and 1.1 compatible)	
On-board Memory	Single-channel with 8 DDR4 8 bit + ECC (64bit + 8bit ECC) **16GB** @ 3200MT/s 64GB eMMC memory	
DPU Power Supply	**Voltage**: 12V	
Power Consumption and Air Flow	Power and airflow specifications are provided in *NVIDIA BlueField-2 DPUs Power and Airflow Specifications* document, which is available at NVOnline following login.	

Environmental	**Temperature**	**Operational**	0°C to 55°C	
		Non-operational	-40°C to 70°C[a]	
	Humidity	**Operational**	10% to 85% relative humidity	
		Non-operational	10% to 90% relative humidity	
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.

BlueField-2 BF2500 Mechanical Drawing and Dimensions

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

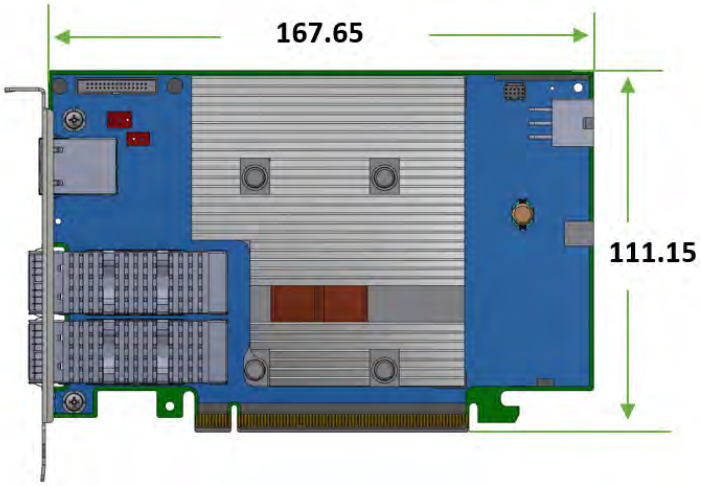

Bracket Mechanical Drawing and Dimensions

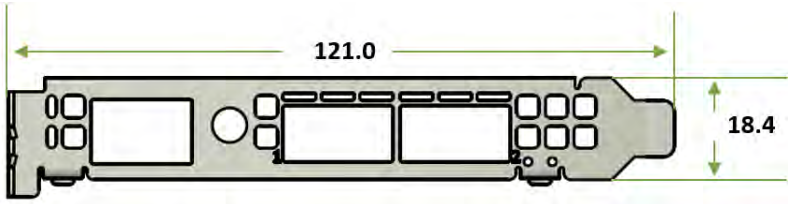

Finding the MAC on the Controller Card

Each DPU Controller has a different identifier printed on the label: serial number and the card MAC for the Ethernet protocol.

> ⚠ The product revisions indicated on the labels in the following figures do not necessarily represent the latest revisions of the cards.

Board Label (Example)

Thermal Sensors

The BlueField-2 BF2500 DPU Controller incorporates the BlueField-2 DPU device which operates in the range of temperatures between 0°C and 105°C.

There are three thermal threshold definitions for the BlueField-2 DPU which impact the overall system operation state:

- **Warning** – 105°C: On managed systems only: When the device crosses the 100°C threshold, a Warning Threshold message will be issued by the management SW, indicating to system administration that the DPU Controller has crossed the Warning threshold. Note that this temperature threshold does not require nor lead to any action by hardware (such as BlueField-2 DPU Controller shutdown).
- **Critical** – 120°C: When the device crosses this temperature, the firmware will automatically shut down the device.
- **Emergency** – 130°C: In case the firmware fails to shut down the device upon crossing the Critical threshold, the device will auto-shutdown upon crossing the Emergency (130°C) threshold.

The BlueField-2 thermal sensors can be read through the system's SMBus. The user can read these thermal sensors and adapt the system airflow in accordance with the readouts and the needs of the above-mentioned DPU thermal requirements.

Document Revision History

Date	Comments/Changes
May. 2023	Updated Specifications - added non-operational storage temperature specifications
Jun. 2021	• Removed OPNs: MBF2H515B-VEEOT, MBF2H515B-VENOT • Multiple text and drawing updates throughout the manual
Dec. 2020	• Updated the NCSI interface pinouts. • Removed the UART interface from "Supported Interfaces". • Updated "Package Contents".
Dec. 2020	Updated the NCSI interface pinouts
Mar. 2020	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.