PRIMERGY
CONVERGED ACCELERATOR CARD
(PY NVIDIA A30X/A100X)
User Manual

1. Introduction

This is the User Manual for CONVERGED ACCELERATOR CARDs (PY NVIDIA A30X/A100X). The cards enable maximum performance on a dense and energy-efficiency form factor when performing artificial intelligence (AI) driven data processing and analytics on edge and datacenter traffic.

Card name	Model Number
PY NVIDIA A30X	PY-CA4A1/PYBCA4A1
PY NVIDIA A100X	PY-CA4A2/PYBCA4A2

Table 1: Product names

1.1. Overview

The cards described in this manual have the following features:

– BlueField-2 DPU with Eight Arm cores and Dual ports 100 Gbps network interface (Ethernet or InfiniBand)

– GPU GA100 tensor core with 24GB (A30X) or 80GB (A100X) memory

– Board Management Controller (BMC) with a RJ45 port

– PCIe interface

  – PCI Express Gen 3.0/4.0 (x16 physical, x8 electrical)

  – Full-Height Full-Length (FHFL), Dual-slot

– NVLink connectors (A30X: 1, A100X: up to 3)

2. CONVERGED ACCELERATOR Card Hardware

2.1. Card covered in this manual

The following table lists the cards described in this manual:

HCA card name	PCI Express	GPU memory	Network I/F	RoHS
PY NVIDIA A30X	PCIe Gen3, 4 (x16 physical, x8 electrical)	24GB	- 100Gbps x2 [QSPF56] (Ethernet/InfiniBand) - 1Gbps x1 [RJ45]	R-10
PY NVIDIA A100X		80GB		R-10

Table 2: PY NVIDIA A30X/A100X

Figure 1: PY NVIDIA A30X

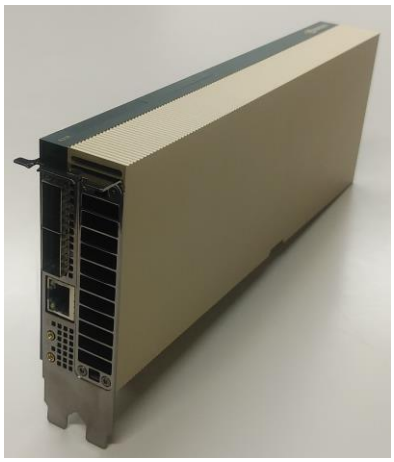

Figure 2: PY NVIDIA A100X

2.2. I/O interface

The cards have the following interfaces:

– Network ports (QSFP56 x2)

– BMC Management LAN port (RJ45 x1)

– PCI Express edge connector (x16 physical, x8 electrical)

– LEDs

2.3. Network port and BMC Management LAN port

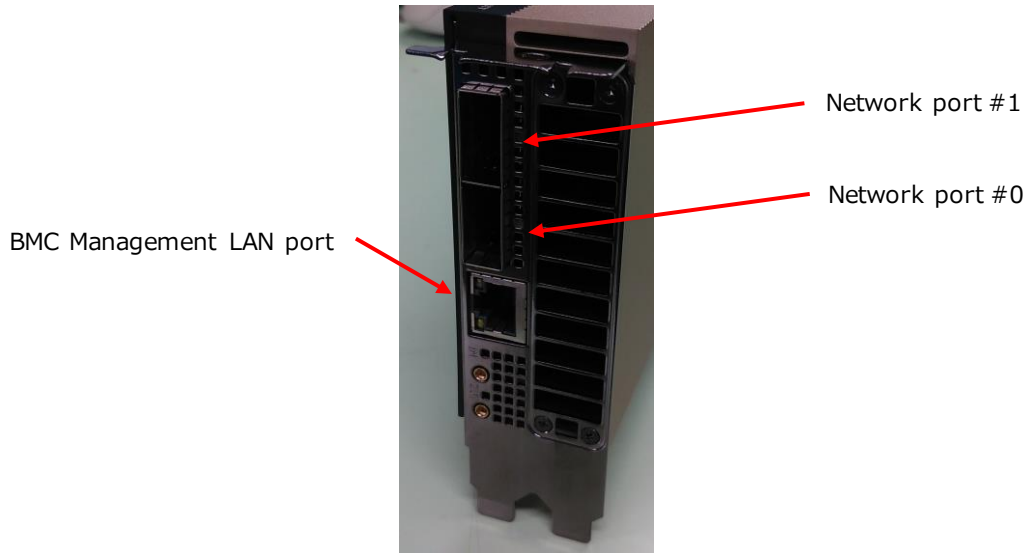

Figure 3: Interface of PY NVIDIA A30X/A100X

CAUTION:

Because default setting of BMC Management LAN port is DHCP mode, connect to the network where the DHCP server is located. After logging into the BMC through the DHCP connection, the static IP address can be configured.

2.3.1. Finding the Serial number, MAC address, and GUID

A card product label on the cards has the Serial number. A DataMatrix (two-dimensional code) label also has the MAC address (Required DataMatrix reader or cell phone with DataMatrix reader application).

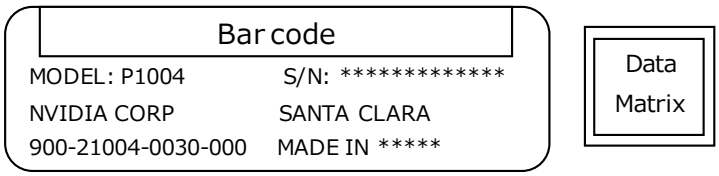

Figure 4: Card product Label and DataMatrix label

The GUID of the network port can be found by running the ibstat commands on the host OS.

CAUTION:

When the card is replaced due to failure, the MAC address and GUID values are changed. Make sure to save those values after building the system environment.

2.3.2. LED assignment

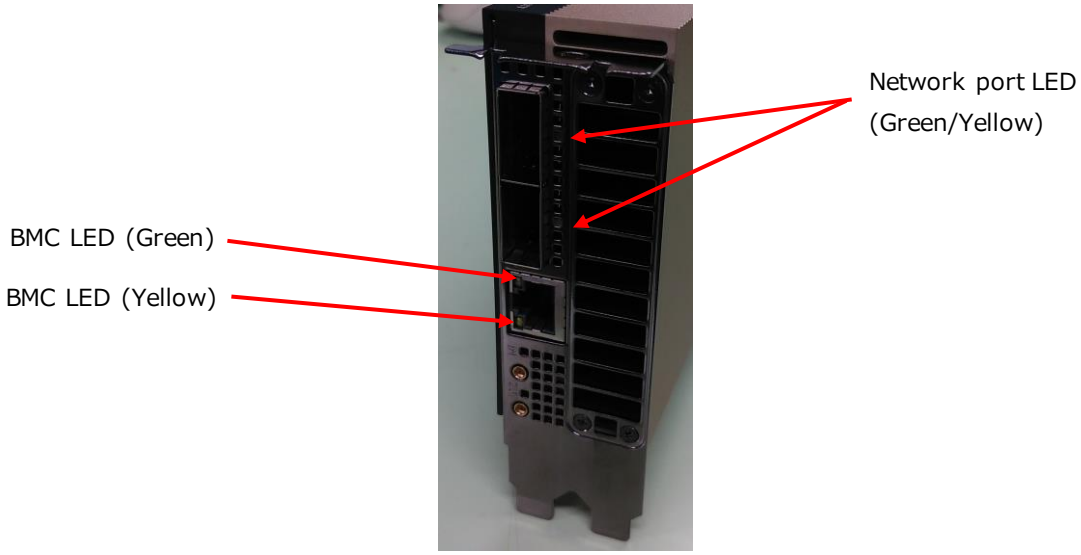

Figure 5: LEDs

Network port LED (bi-color)	Status	Function
-	OFF	Not link up
Yellow	BLINK	Error on link
Green	ON	Link up
	BLINK	In data traffic

BMC LED		Function
Green	Yellow	
OFF	OFF	Not link up
ON	OFF	Link up
ON	BLINK	In data traffic

Table 3: LED assignment

2.4. PCI Express interface

The cards support PCI Express 3.0 and 4.0 (1.1, 2.0 compatible) through an edge connector (x16 physical, x8 electrical). The cards can be either a master initiating the PCI Express bus operations or a slave responding to PCI bus operations.

3. CONVERGED ACCELERATOR CARD installation

3.1. System requirement and Attention

For the system requirements and attentions on mounting the cards on the server, refer to "FUJITSU WebArchitect – Product & System Configurator".
https://www.fujitsu.com/emeia/products/configurator/

CAUTION:

When a card is replaced due to a failure, the DPU and BMC settings are initialized and the MAC address or GUID values are changed. Make sure to save settings and those values after building the system environment and reconfigure them if necessary.

3.2. Installation procedure

For the installation procedure into the servers, refer to "Upgrade and Maintenance Manual" for each server.
https://support.ts.fujitsu.com/IndexDownload.asp

CAUTION:

When installing or removing the card, turn off the power to each device (Servers, peripherals, etc.) and disconnect the power cord from the wall outlet. If you work with the power cable attached, you may get an electric shock. After installing the card, connect the power cord to the server.

4. Driver installation and Firmware update

4.1. Driver download

The driver for the cards can be found at the following website. Download the appropriate version of the driver if necessary.

・NVIDIA GPU driver

Select Data Center/Tesla, A-Series, NVIDIA A100.

https://www.nvidia.co.jp/Download/index.aspx?lang=en

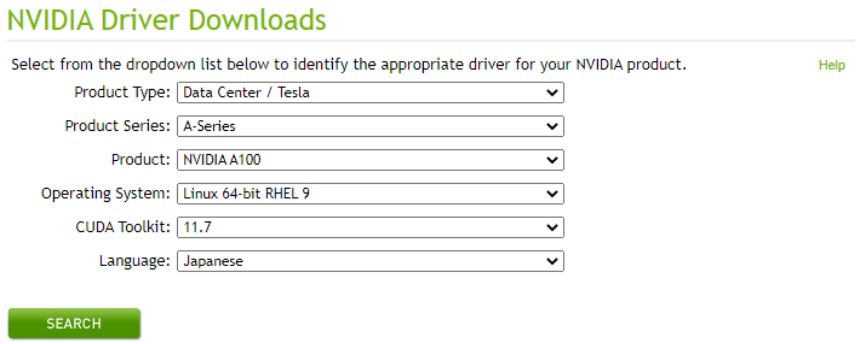

Figure 6: GPU driver download webpage

・NVIDIA MLNX_OFED driver

https://network.nvidia.com/products/infiniband-drivers/linux/mlnx_ofed/

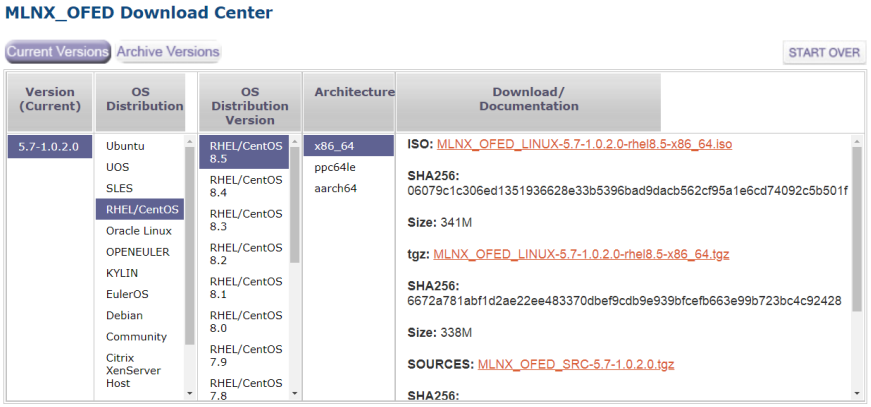

Figure 7: OFED driver download webpage

4.2. Driver installation and uninstallation

The procedures for installing and uninstalling the driver can be found in the “Release Notes” that appear when you select the download driver. Additional packages installation may be required during installation. Follow the message to install the packages.

・NVIDIA GPU driver Release Notes

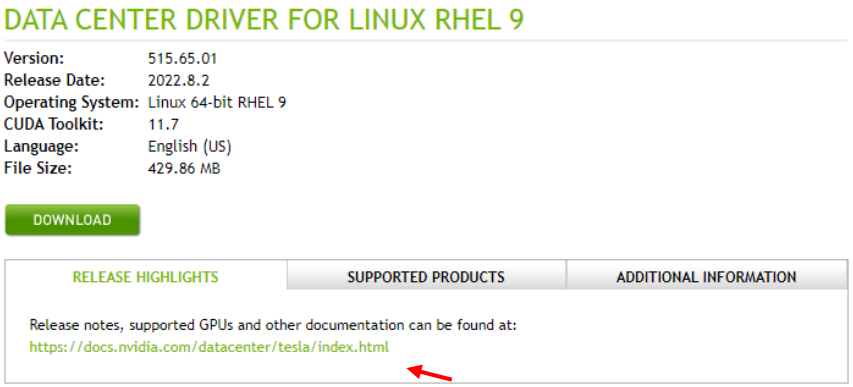

Figure 8: GPU driver Release Notes download webpage

・NVIDIA MLNX_OFED driver Release Notes

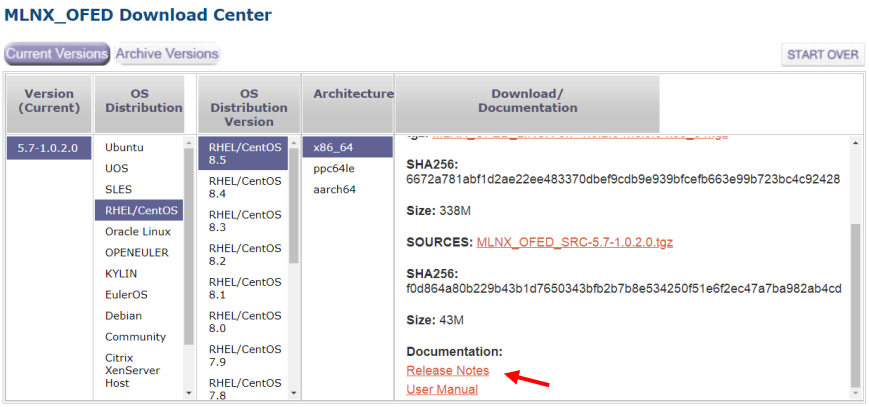

Figure 9: OFED driver Release Notes download webpage

CAUTION:

Do not apply the driver attached to the OS when updating the OS. Apply the driver downloaded from the website above.

4.3. Downloading Firmware and BlueField Bootstream (BFB)

The Firmware for the cards can be found at the following website. Download the appropriate version of the driver if necessary.

・NVIDIA BlueField-2 DPU Firmware

https://network.nvidia.com/support/firmware/bluefield2/

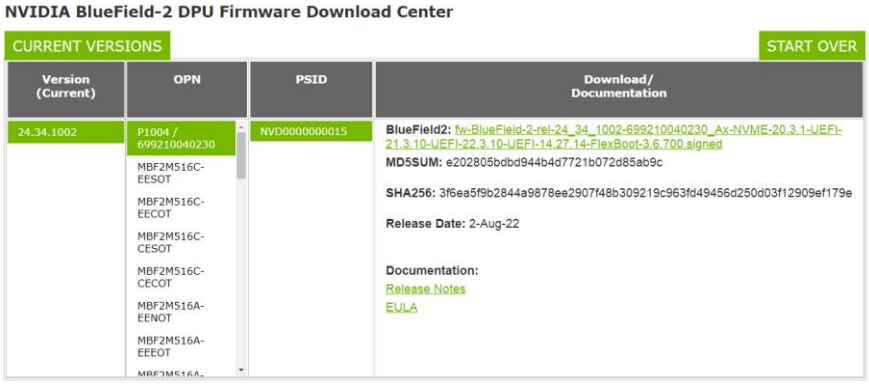

Figure 10: DPU Firmware download webpage

・BMC and CEC Firmware

https://developer.nvidia.com/networking/doca

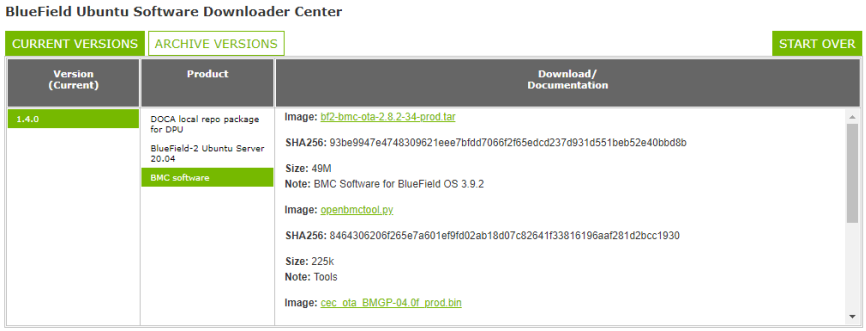

Figure 11: BMC/CEC Firmware download webpage

・BlueField Bootstream (BFB)

https://developer.nvidia.com/networking/doca

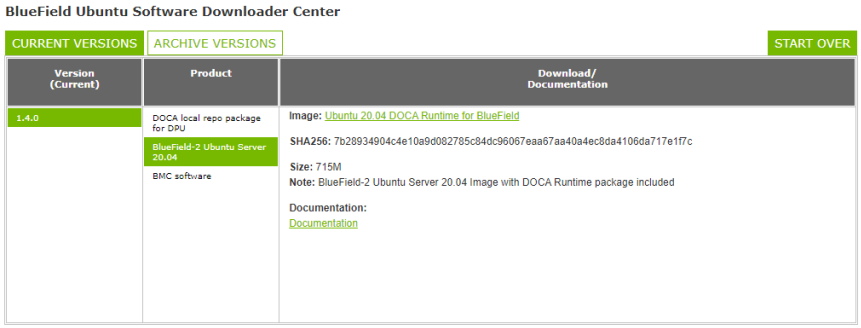

Figure 12: BMC/CEC Firmware download webpage

4.4. Applying Firmware and BlueField Bootstream (BFB)

・NVIDIA BlueField-2 DPU, BMC, and CEC Firmware ,

The procedures for applying the downloaded firmware can be found on the following website:

https://docs.nvidia.com/networking/category/bluefieldsw

Figure 13: BMC/CEC Firmware documents download webpage

・BlueField Bootstream (BFB)

The procedure for applying the downloaded BFB (DPU OS) can be found on the following website:

https://docs.nvidia.com/doca/sdk/installation-guide/index.html

5. Replacing CONVERGED ACCELERATOR Card

CAUTION:

The precautions when trouble occurs are following:

1. If user encounter any problems using an Independent Software Vendor (ISV) DPU application or has any questions about the application, be sure to first contact the ISV for assistance. If the ISV reports that there is no problem with the DPU application and there is a problem with the card, contact our company support window. Our company cannot handle problems and inquiries with the DPU application.
2. Our company support team may ask user to submit the survey results of ISV for investigation.
3. We may ask user to provide the system event log and dump of the DPU for investigation. In this case, the user must collect them in accordance with Section 6.
4. User may not be able to save and export the setting values and user data to other location due to failure. After the building system environment has been done and the settings have been changed, be sure to store the data and record the settings to other location.

CAUTION:

The precautions when replacing the card are following;

1. Use an earth strap to prevent electrostatic damage. Also, be sure to place the replacement card on an electrostatic mat.
2. The firmware must be installed by the customer in accordance with the requirements of the DPU application developed by ISV. If there is no firmware version requirement, the latest version is recommended. It is also recommended that all firmware versions be aligned when multiple these cards are used in a server and there are no DPU application firmware requirements. Refer to section 4.3 for downloading the firmware and BFB
3. The MAC address and GUID change when the card is replaced. After replacing the card, the user must confirm the new MAC address and GUID in accordance with Section 2.3.1.
4. BFB and BMC settings and user data are initialized by card replacement. Before

replacing the card, the user must save the BFB and BMC settings and user data to other location. After replacing the card, the user must reconfigure the saved settings and user data. Refer to the documentations in Section 6 for how to operate the DPU and BMC.

5. After replacing the card, the BMC management LAN IP address may be changed by the DHCP server. If the Static IP address has been set, the user must reconfigure the Static IP address according to the DPU/BMC vendor documentation in Section 6.

5.1. Card Replacement Procedures

① Checking the cabling location
Check the card slot location and module and cable location for reconfiguration after card replacement.

② Collection of card information and fault information
Obtain the card information (such as FW and BFB revisions) and fault information, and retrieve the DPU system event log and dump. Our company support team may ask user to provide them. Refer to DPU/BMC vendor documentations in Section 6 how to operate the DPU and BMC.

③ Saving setting values and user data
After replacing the card, the settings and user data are initialized, so the user must save them to other location before replacing the card.

④ Removing a fault card
Remove all modules and cables from the fault card and the card from the server. Refer to the "Upgrade and Maintenance Manual" in Section 3.2 for removal.

⑤ Checking new MAC address of new card
Users must check the new MAC address with a reader that can read the DataMatrix label on the new card or a cell phone with the reader application installed. Be sure to record and store the scanned MAC address. Refer to Section 2.3.1 for label information.

⑥ Installing new card
Install the new card in the same slot as before, and then install the modules and cables in the card. Refer to the “Upgrade and Maintenance Manual” in Section 3.2 for installation.

⑦ Checking new GUID of new card

After the server boot, the user must check for the new GUID. The GUID can be checked by running the ibstat command on the host OS.

⑧ Firmware update (may be before or after ⑨)
Update to the same firmware version as before. Refer to Section 4.7 for the firmware update.

⑨ Reconfiguring settings and user data (may be before or after ⑧)
The user must reconfigure all the contents saved in ③.

6. The latest vendor product document

The latest documents for NVIDIA A100X/A30X can be found on the following website:

・NVIDIA CONVERGED ACCELERATOR introduction

https://www.nvidia.com/ja-jp/data-center/products/converged-accelerator/

・DPU/BMC (BlueField Family Products)

https://docs.nvidia.com/networking/category/bluefieldsw

・Network (Ethernet/InfiniBand)

https://docs.nvidia.com/networking/category/bluefieldsnic

・NVIDIA DOCA Software Framework

https://developer.nvidia.com/networking/doca

7. Specifications

PY-CA4A1/PYBCA4A1

Physical	**[Size]** 266.9mm x 98.1mm x 38.5mm Full-Height Full-Length(FHFL), Dual-slot
	[Connector] ・QSFP56 (100Gbps Ethernet or InfiniBand) ・RJ-45 (1Gbps Ethernet, Required DHCP Server) ・NVLink ・Power cable (8 pin)
Protocol Support	**[Ethernet]** InfiniBand: IBTA v1.3 **[InfiniBand]** - 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)
	[Data Rate] Up to 100Gbps (Ethernet) Up to 100Gb/s EDR (InfiniBand)
	[PCI Express] PCI Express Gen3 and 4 (2.0 and 1.1 compatible), x16 physical / x8 electrical

Power and Environmental	**[Voltage]** 12V, 3.3V (PCIe card edge) 12V (Power cable)
	[Power] Max 230W
	[Temperature] 0°C to 50°C
Regulatory	**[Safety]** CB/UL **[EMC]** RCM, CE, UKCA, VCCI, KC, Morocco, SABS, BSMI, Ukraine, FCCCI/ICES
	[RoHS] R-10

Table 4: PY-CA4A1/PYBCA4A1 specification

PY-CA4A2/PYBCA4A2

Physical	**[Size]** 266.9mm x 98.1mm x 38.5mm Full-Height Full-Length(FHFL), Dual-slot
	[Connector] ・QSFP56 (100Gbps Ethernet or InfiniBand) ・RJ-45 (1Gbps Ethernet, Required DHCP Server) ・NVLink (up to 3) ・Power cable（8 pin）
Protocol Support	**[Ethernet]** InfiniBand: IBTA v1.3 **[InfiniBand]** - 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)
	[Data Rate] Up to 100Gbps (Ethernet) Up to 100Gb/s EDR (InfiniBand)
	[PCI Express] PCI Express Gen3 and 4 (2.0 and 1.1 compatible), x16 physical / x8 electrical
Power and Environmental	**[Voltage]** 12V, 3.3V (PCIe card edge) 12V (Power cable)

	[Power] Max 300W
	[Temperature] 0°C to 50°C
Regulatory	**[Safety]** CB/UL **[EMC]** RCM, CE, UKCA, VCCI, KC, Morocco, SABS, BSMI, Ukraine, FCCCI/ICES
	[RoHS] R-10

Table 5: PY-CA4A2/PYBCA4A2 specification

PRIMERGY COVERGED ACCELERATOR CARD (PY NVIDIA A30X/A100X)

User Manual

CA92344-5055 ver. 1.0

Issued on Sep, 2022

Issued by FUJITSU LIMITED

Printed in Japan

- The contents may be revised without prior notice.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- No part of this manual may be reproduced in any form without the prior written permission of Fujitsu.
- Manuals with missing or wrongly collated pages will be replaced free of charge Copyright and Trademarks