OCP Mezzanine Card 2.0 Product Specification

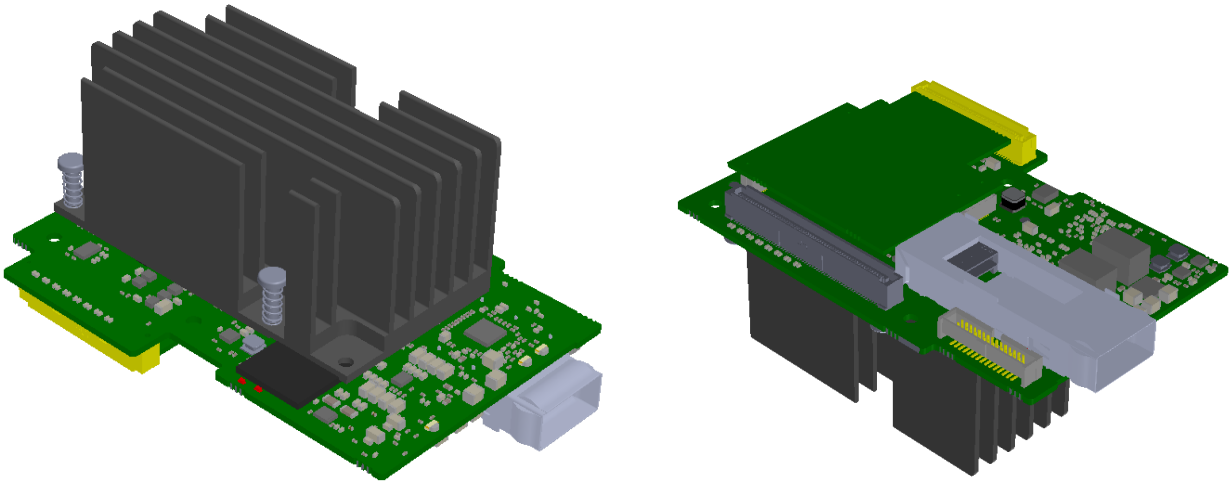

Netronome's 50G Crypto Mezz 2.0 Card

Version 1.3

Author:

Uwe PACK, Principal Engineer, Netronome Systems, Inc.

Legal Disclaimer

As of May 20, 2019, the following persons or entities have made this Specification available under the Open Web Foundation Contributors License Agreement 1.0 –Copyright and Patent, which is available at:

http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owf-contributor-license-agreement-1-0---copyright-and-patent

Netronome Systems, Inc.

Your use of this Specification may be subject to other third party rights. THIS SPECIFICATION IS PROVIDED "AS IS." Netronome and its contributors expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, noninfringement, fitness for a particular purpose, or title, related to the Specification. The entire risk as to implementing or otherwise using the Specification is assumed by the Specification implementer and user. IN NO EVENT WILL ANY PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CONTRIBUTORS AND LICENSORS OF THIS SPECIFICATION MAY HAVE MENTIONED CERTAIN TECHNOLOGIES THAT ARE MERELY REFERENCED WITHIN THIS SPECIFICATION AND NOT LICENSED UNDER THE OWF CLA OR OWFa. IMPLEMENTATION OF THESE TECHNOLOGIES, WHICH MAY INCLUDE SAMTEC TECHNOLOGY, MAY BE SUBJECT TO THEIR OWN LEGAL TERMS.

Version History

Date	Description	Version
03/8/2019	Initial Draft version 1.0	Draft 1.0
03/12/2019	Correction in Overview section	Draft 1.1
05/17/2019	Added Power and Mechanical details	1.2
05/20/2019	Header & Footer correction	1.3

Table of Contents

Table of Tables

Tables of Figures

1. Overview

This document defines physical and interface requirements for Netronome OCP Mezzanine 2.0 cards that can be installed in an OCP Multi-Host Server, such as Yosemite-v2.

These OCP Mezzanine 2.0 cards are programmable NICs that provide CPU offload for Host-based SDN, virtual switch data path, and tunneling protocols.

Described is a 50G network connection using a single QSFP-28 interface with Crypto enabled.

This specification does not cover the functionality of the programmable Netronome OCP Mezzanine cards that will be developed per these physical requirements.

1.1. Block Diagram

Netronome OCP Mezzanine cards are installed on the OCP Multi-Host Servers via PCIe Gen3 interface through the baseboard to the central processing unit on each server. Block diagrams of the essential elements of the card is shown below:

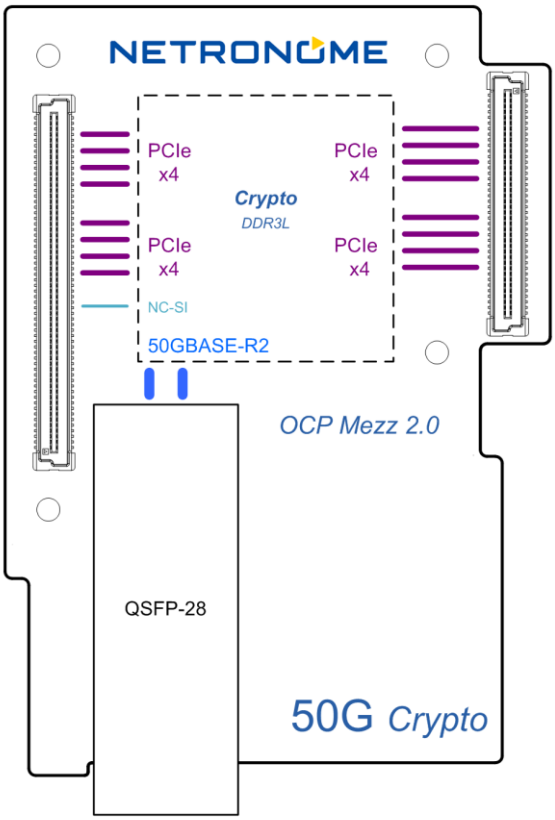

Figure 1: Netronome OCP Mezzanine Card Block Diagram

2. Connectors and Pinouts

2.1. Board Connector

The Netronome OCP Mezzanine card connects to the motherboard using FCI / Amphenol's BergStak 0.8mm Mezzanine connector solution.
For completeness, the table below lists both Mezzanine and Baseboard connectors for various configurations, with part numbers not applicable to these cards grayed out.

Table 1: Connector Part Numbers

	Mezzanine card (5mm stack)	Mezzanine card (8/12mm stack)	Baseboard (5/8mm stack)	Baseboard (12mm stack)
Connector A	FCI/61083-121402LF	FCI/61083-124402LF	FCI/61082-121402LF	FCI/61082-122402LF
Connector B	FCI/61083-081402LF	FCI/61083-084402LF	FCI/61082-081402LF	FCI/61082-082402LF
Connector C	FCI/10135584-641402LF	FCI/10135584-644402LF	FCI/10135583-641402LF	FCI/10135583-642402LF

The Netronome OCP Mezzanine cards only use Connector-A and Connector-B (120 and 80 positions, respectively) to interface with the baseboard.

2.2. Signal Definitions

Table 2 details the signals used in the Mezzanine-baseboard interface on both Netronome OCP cards.

Table 2: Mezzanine Card Pin Description

Signals on Connector A	*Type*	*Description*
GND	Ground	Ground return; total 51 pins on Connector A
P12V_AUX/P12V	Power	12V Aux/normal power; total 3 pins on Connector A
P5V_AUX	Power	5V Aux power; total 3 pins on Connector A
P3V3_AUX	Power	P3V3 Aux Power; total 2 pins on Connector A
P3V3	Power	P3V3 power; total 8 pins on Connector A
MEZZ_PRSNTA1_N/BASEBOARD_ID_A	Output	Connector A Present Pin; connect to MEZZ_PRSNTA2_N on Mezz with 0 Ohm; Use as baseboard ID during power up
MEZZ_PRSNTA2_N	Input	Connector A Present Pin; connect to MEZZ_PRSNTA1_N on Mezz with 0 Ohm
LAN_3V3STB_ALERT_N	Input	SMBus Alert for OOB management; 3.3V AUX rail
SMB_LAN_3V3STB_CLK	Output	SMBus Clock for OOB management; 3.3V AUX rail; Share with thermal reporting interface; Both 100Kb/s and 400Kb/s shall be supported
SMB_LAN_3V3STB_DAT	Bidirectional	SMBus Data for OOB management; 3.3V AUX rail; Share with thermal reporting interface; Both 100Kb/s and 400Kb/s shall be supported

NCSI_RXER	Input	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_CRSDV	Input	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_RXD[1..0]	Input	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_RCLK	Output	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_TXEN	Output	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_TXD[1..0]	Output	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
PCIE_WAKE_N	Input	PCIe wake up signal
PERST_N0	Output	PCIe reset signal 0
MEZZ_SMCLK	Output	PCIe SMBus Clock for Mezz slot/EEPROM; 3.3V AUX rail; Share with thermal reporting interface; Both 100Kb/s and 400Kb/s shall be supported
MEZZ_SMDATA	Bidirectional	PCIe SMBus Data for Mezz slot/EEPROM; 3.3V AUX rail; Share with thermal reporting interface; Both 100Kb/s and 400Kb/s shall be supported
CLK_100M_MEZZ[1..0]_DP/N	Output	MB clock output for PCIe devices; total 2 pairson Connector A; CLK_100M_MEZZ1_DP/N is optional for single host baseboard
NCSI_RXER	Input	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
NCSI_CRSDV	Input	NC‐SI for OOB management; 3.3V AUX rail; Direction is in perspective of baseboard
MEZZ_TX_DP/N_C<7..0>	Output	PCIe TX; total up to 8 lanes on Connector A; optional with KR signals
MEZZ_RX_DP/N<7..0>	Input	PCIe RX; total up to 8 lanes on Connector A; optional with KR signals
KR_TX_DP/N<15..8>	Output	KR TX; total up to 8 lanes on Connector A; optional with PCIe signals
KR_RX_DP/N<15..8>	Input	KR RX; total up to 8 lanes on Connector A; optional with PCIe signals
RSVD	TBD	Reserved for Future use

Signals on Connector B	*Type*	*Description*
GND	Ground	Ground return; total 36 pins on Connector B
P12V_AUX/P12V	Power	12V Aux/Normal power; total 2 pins on Connector B
MEZZ_PRSNTB1_N/ BASEBOARD_ID_B	Output	Connector B Present Pin; connect to MEZZ_PRSNTB2_N on Mezz with 0 Ohm Use as baseboard ID during power up
MEZZ_PRSNTB2_N	Input	Connector B Present Pin; connect to MEZZ_PRSNTB1_N on Mezz with 0 Ohm
PERST_N[3..1]	Output	PCIe reset signal or Node[3..1] PCIe reset signal for baseboard with more than 1 nodes
CLK_100M_MEZZ[3..2]_DP/N	Output	MB clock output for PCIe devices; total 2 pairs on Connector B; optional for single host baseboard
MEZZ_TX_DP/N_C<15..8>	Output	PCIe TX; total up to 8 lanes on Connector B; optional with KR signals
MEZZ_RX_DP/N<15..8>	Input	PCIe RX; total up to 8 lanes on Connector B; optional with KR signals
KR_TX_DP/N<7..0>	Output	KR TX; total up to 8 lanes on Connector B; optional with PCIe signals
KR_RX_DP/N<7..0>	Input	KR RX; total up to 8 lanes on Connector B; optional with PCIe signals
RSVD	TBD	Reserved for Future use

2.3. Connector Pinout

Table 3: x16 PCIe Mezzanine Card Pin Definition

Connector A

Signal	Pin	Pin	Signal
P12V_AUX/P12V	A61	A1	MEZZ_PRSNTA1_N/ BASEBOARD_A_ID
P12V_AUX/P12V	A62	A2	P5V_AUX
P12V_AUX/P12V	A63	A3	P5V_AUX
GND	A64	A4	P5V_AUX
GND	A65	A5	GND
P3V3_AUX	A66	A6	GND
GND	A67	A7	P3V3_AUX
GND	A68	A8	GND
P3V3	A69	A9	GND
P3V3	A70	A10	P3V3
P3V3	A71	A11	P3V3
P3V3	A72	A12	P3V3
GND	A73	A13	P3V3
LAN_3V3STB_ALERT_N	A74	A14	NCSI_CRSDV
SMB_LAN_3V3STB_CLK	A75	A15	NCSI_RCLK
SMB_LAN_3V3STB_DAT	A76	A16	NCSI_TXEN
PCIE_WAKE_N	A77	A17	PERST_N0
NCSI_RXER	A78	A18	MEZZ_SMCLK
GND	A79	A19	MEZZ_SMDATA
NCSI_TXD0	A80	A20	GND
NCSI_TXD1	A81	A21	GND
GND	A82	A22	NCSI_RXD0
GND	A83	A23	NCSI_RXD1
CLK_100M_MEZZ0_DP	A84	A24	GND
CLK_100M_MEZZ0_DN	A85	A25	GND
GND	A86	A26	CLK_100M_MEZZ1_DP
GND	A87	A27	CLK_100M_MEZZ1_DN
MEZZ_TX_DP_C<0>	A88	A28	GND
MEZZ_TX_DN_C<0>	A89	A29	GND
GND	A90	A30	MEZZ_RX_DP<0>
GND	A91	A31	MEZZ_RX_DN<0>
MEZZ_TX_DP_C<1>	A92	A32	GND
MEZZ_TX_DN_C<1>	A93	A33	GND
GND	A94	A34	MEZZ_RX_DP<1>
GND	A95	A35	MEZZ_RX_DN<1>
MEZZ_TX_DP_C<2>	A96	A36	GND
MEZZ_TX_DN_C<2>	A97	A37	GND
GND	A98	A38	MEZZ_RX_DP<2>
GND	A99	A39	MEZZ_RX_DN<2>
MEZZ_TX_DP_C<3>	A100	A40	GND
MEZZ_TX_DN_C<3>	A101	A41	GND
GND	A102	A42	MEZZ_RX_DP<3>
GND	A103	A43	MEZZ_RX_DN<3>
MEZZ_TX_DP_C<4>	A104	A44	GND
MEZZ_TX_DN_C<4>	A105	A45	GND
GND	A106	A46	MEZZ_RX_DP<4>
GND	A107	A47	MEZZ_RX_DN<4>
MEZZ_TX_DP_C<5>	A108	A48	GND
MEZZ_TX_DN_C<5>	A109	A49	GND

Connector B

Signal	Pin	Pin	Signal
P12V_AUX/P12V	B41	B1	MEZZ_PRSNTB1_N/ BASEBOARD_B_ID
P12V_AUX/P12V	B42	B2	GND
RSVD	B43	B3	MEZZ_RX_DP<8>
GND	B44	B4	MEZZ_RX_DN<8>
MEZZ_TX_DP<8>	B45	B5	GND
MEZZ_TX_DN<8>	B46	B6	GND
GND	B47	B7	MEZZ_RX_DP<9>
GND	B48	B8	MEZZ_RX_DN<9>
MEZZ_TX_DP<9>	B49	B9	GND
MEZZ_TX_DN<9>	B50	B10	GND
GND	B51	B11	MEZZ_RX_DP<10>
GND	B52	B12	MEZZ_RX_DN<10>
MEZZ_TX_DP<10>	B53	B13	GND
MEZZ_TX_DN<10>	B54	B14	GND
GND	B55	B15	MEZZ_RX_DP<11>
GND	B56	B16	MEZZ_RX_DN<11>
MEZZ_TX_DP<11>	B57	B17	GND
MEZZ_TX_DN<11>	B58	B18	GND
GND	B59	B19	MEZZ_RX_DP<12>
GND	B60	B20	MEZZ_RX_DN<12>
MEZZ_TX_DP<12>	B61	B21	GND
MEZZ_TX_DN<12>	B62	B22	GND
GND	B63	B23	MEZZ_RX_DP<13>
GND	B64	B24	MEZZ_RX_DN<13>
MEZZ_TX_DP<13>	B65	B25	GND
MEZZ_TX_DN<13>	B66	B26	GND
GND	B67	B27	MEZZ_RX_DP<14>
GND	B68	B28	MEZZ_RX_DN<14>
MEZZ_TX_DP<14>	B69	B29	GND
MEZZ_TX_DN<14>	B70	B30	GND
GND	B71	B31	MEZZ_RX_DP<15>
GND	B72	B32	MEZZ_RX_DN<15>
MEZZ_TX_DP<15>	B73	B33	GND
MEZZ_TX_DN<15>	B74	B34	GND
GND	B75	B35	CLK_100M_MEZZ2_DP
GND	B76	B36	CLK_100M_MEZZ2_DN
CLK_100M_MEZZ3_DP	B77	B37	GND
CLK_100M_MEZZ3_DN	B78	B38	PERST_N1
GND	B79	B39	PERST_N2
MEZZ_PRSNTB2_N	B80	B40	PERST_N3

GND	A110	A50	MEZZ_RX_DP<5>
GND	A111	A51	MEZZ_RX_DN<5>
MEZZ_TX_DP_C<6>	A112	A52	GND
MEZZ_TX_DN_C<6>	A113	A53	GND
GND	A114	A54	MEZZ_RX_DP<6>
GND	A115	A55	MEZZ_RX_DN<6>
MEZZ_TX_DP_C<7>	A116	A56	GND
MEZZ_TX_DN_C<7>	A117	A57	GND
GND	A118	A58	MEZZ_RX_DP<7>
GND	A119	A59	MEZZ_RX_DN<7>
MEZZ_PRSNTA2_N	A120	A60	GND

3. Power Consumption

The following tables detail the power ratings for each of the power rails on each connector.

Table 4: Power Ratings for Connector A

Power Rail	Voltage Tolerance	# of pins	Current Capability	Status
P12V_AUX/P 12V	±8%(max)	3	2.4A	Auxiliary Power/Normal Power
P5V_AUX	±9%(max)	3	2.4A	Auxiliary power
P3V3_AUX	±5%(max)	2	1.6A	Auxiliary power
P3V3	±5%(max)	8	6.4A	Normal power

Table 5: Power Ratings for Connector B

Power Rail	Voltage Tolerance	# of pins	Current Capability	Status
P12V_AUX/P 12V	±8%(max)	2	1.6A	Auxiliary Power/Normal Power

Netronome's OCP Mezzanine card only draws power from Connector A, and each of the AUX rails is monitored for power consumption.

Table 6: Netronome's Maximum Power per Rail

Power Rail	Current Capability	Worst-Case Current	Status
P12V_AUX/P 12V	2.4A (Conn-A *only*)	< 2.2A	Power Monitored
P5V_AUX	2.4A	< 1.6A	Power Monitored
P3V3_AUX	1.6A	< 0.8A	Power Monitored
P3V3	6.4A	~ 10mA	Presence

Power form each rail is added and software will 'throttle' performance to remain under 25W total power consumption.

4. Mechanical

The form-factor of this Mezzanine card is closest to a *Type-1* card with our heat sink violating the *Secondary Side* height restrictions.

4.1. Horizontal Plane (X and Y)

The following figure details the outline horizontal plane view and dimensions for Netronome's OCP Mezzanine card.

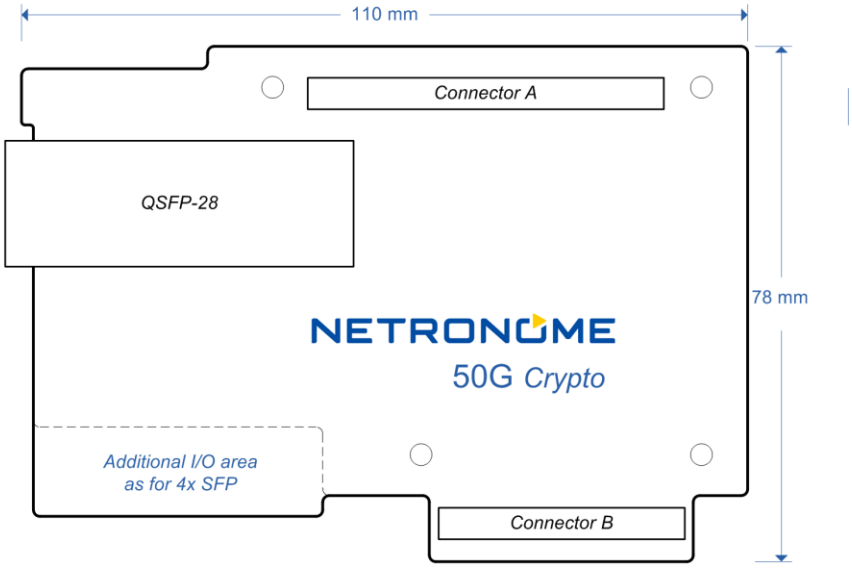

Figure 2: Mechanical Layout of Netronome OCP Mezzanine Card

4.2 Vertical (Z-axis)

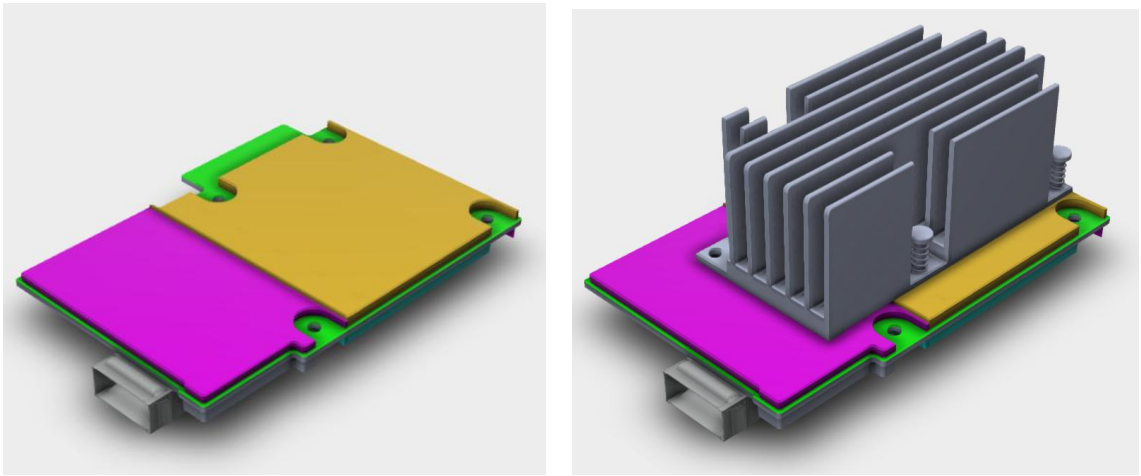

Figure 3: Heat sink violates Secondary Side height restrictions

For Yosemite-v2, the heat sink is 38mm tall.

Below are different heat sink options for *Type-1* or *Type-2* for a 1U server.

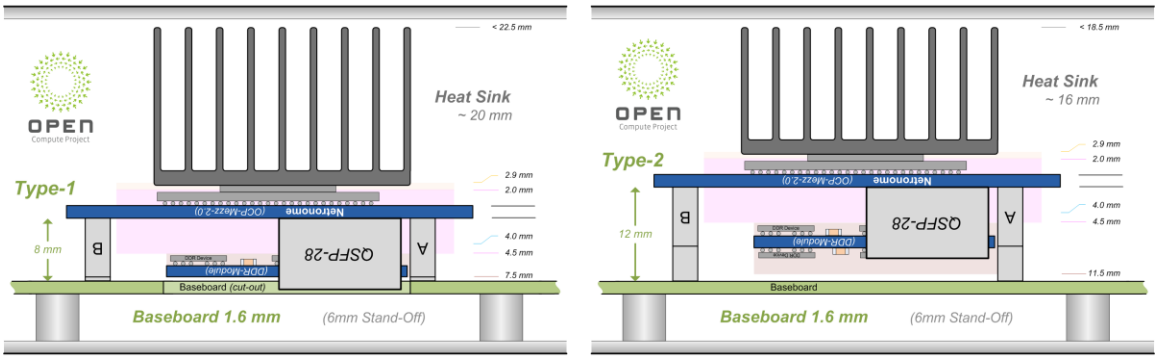

Figure 4: Front- View of 50G Netronome Mezzanine Card

This Mezzanine has a DDR-Module which may interfere with component on the base board, especially for a *Type-1* application.

- **7.5mm height restriction** is violated by 0.1mm – may be okay
- **4.5mm height restriction** and **4.0mm height restriction** would interfere, if components are placed in that region

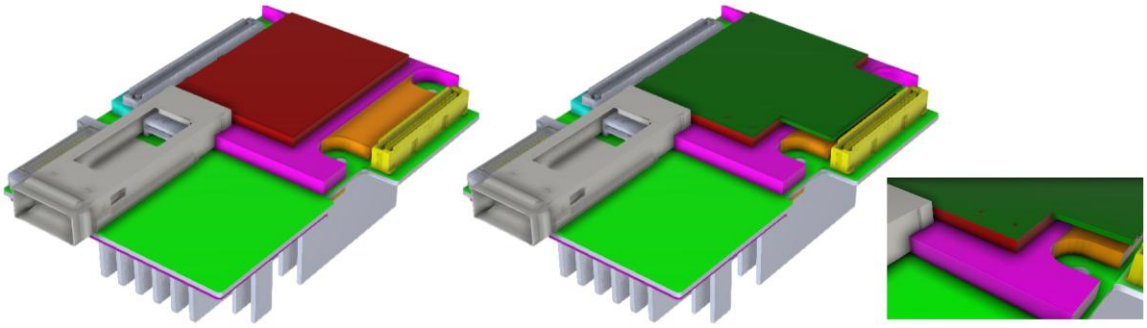

Figure 3: DDR-Module interference on Primary Side

5. Thermal Requirements

The heat sink needs to be application specific for the environment. A taller heat sink with fewer fins is more efficient than a shallow heat sink with many fins in low airflow environment.

For Yosemite-v2, the heat sink is optimized for relative low airflow (about 1m/s) at 30°C, and requires >2.25m/s (450LFM) airflow at 45°C inlet temperature.