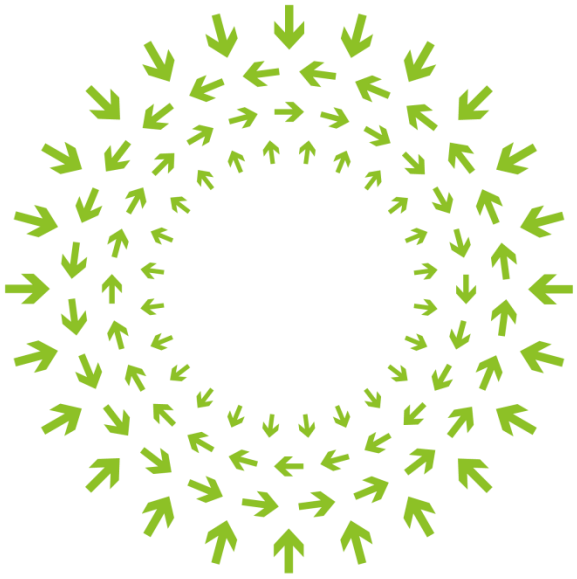

OPEN

Compute Project

Open CloudServer OCS
NIC Mezzanine Specification
Version 2.0

Author:
Mark Shaw, Director of Hardware Engineering, Microsoft

Revision History

Date	Description
10/10/2014	Version 2.0

Contents

Table of Figures

Table of Tables

1 Summary

This specification, *Open CloudServer NIC Mezzanine Version 2.0*, describes the physical and interface requirements for the Open CloudServer (OCS) NIC mezzanine card that to be installed on an OCS blade.

Note that this specification does not cover the functional or features requirements for the mezzanine cards that will be developed for this project.

1.1 Block Diagram

The mezzanine card will be installed on the blade and will have a Peripheral Component Interconnect Express (PCIe) x8 Gen3 interface to a central processing unit (CPU) on the blade motherboard. It will support either a single 10Gigabit Ethernet (GbE) to a small form-factor pluggable (SFP)+ cable or a single 40GbE interface to a quad small form-factor pluggable (QSFP)+ cable. It is never required to support both interfaces on the same mezzanine card.

To maintain compatibility with existing mezzanine designs, the 10GbE port must map to network port 1 on the SEAF/SEAM connector interface, and the 40GbE port must map to network port 2 on the SEAF/SEAM connector interface. Figure 1 shows a block diagram of the interface.

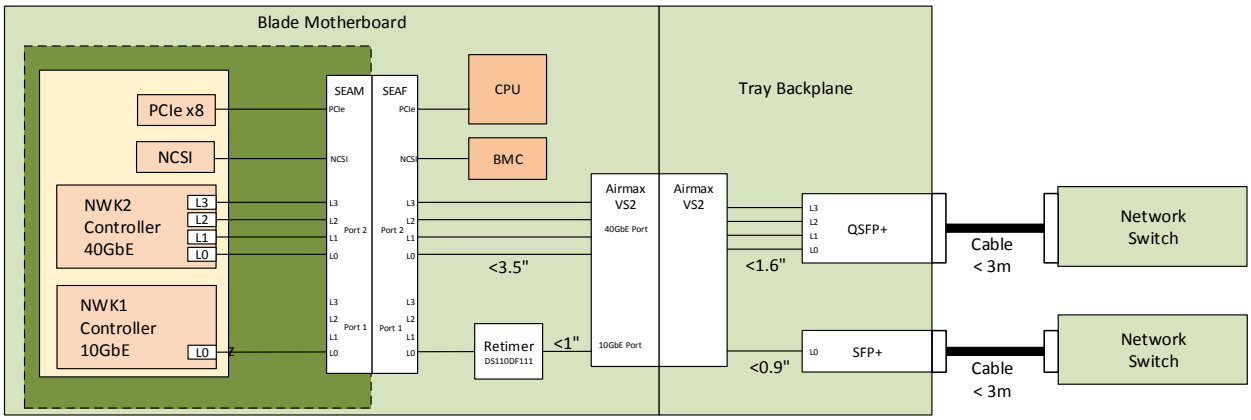

Figure 1: NIC mezzanine block diagram

1.2 Interface Topology

Figure 1 above shows the full topology of the GbE interface, including maximum link lengths. Detailed design information for the motherboard and tray backplane is available upon request.

2 Interface Connector

The connector interfaces between the NIC mezzanine card and the motherboard use the Samtec SEARAY solution. The mezzanine card uses a Samtec SEAM (male) connector and interfaces to the motherboard through a SEAF (female) connector. Table 1 lists the manufacturing part numbers (MPNs).

Table 1: Connector manufacturing part numbers

Manufacturer	Mezzanine card connector MPN	Mating (MB) connector MPN
Samtec	SEAM-20-03.5-S-08-2-A-K-TR	SEAF-20-06.5-S-08-2-A-K-TR
Molex	45970-2385	45971-2385

The stackup height of the SEAM is 10mm, with a 6.5mm SEAF and a 3.5mm SEAM connector. In this configuration, it is expected that taller components would be place on the top side of the printed circuit board (PCB), as shown below in Figure 2.

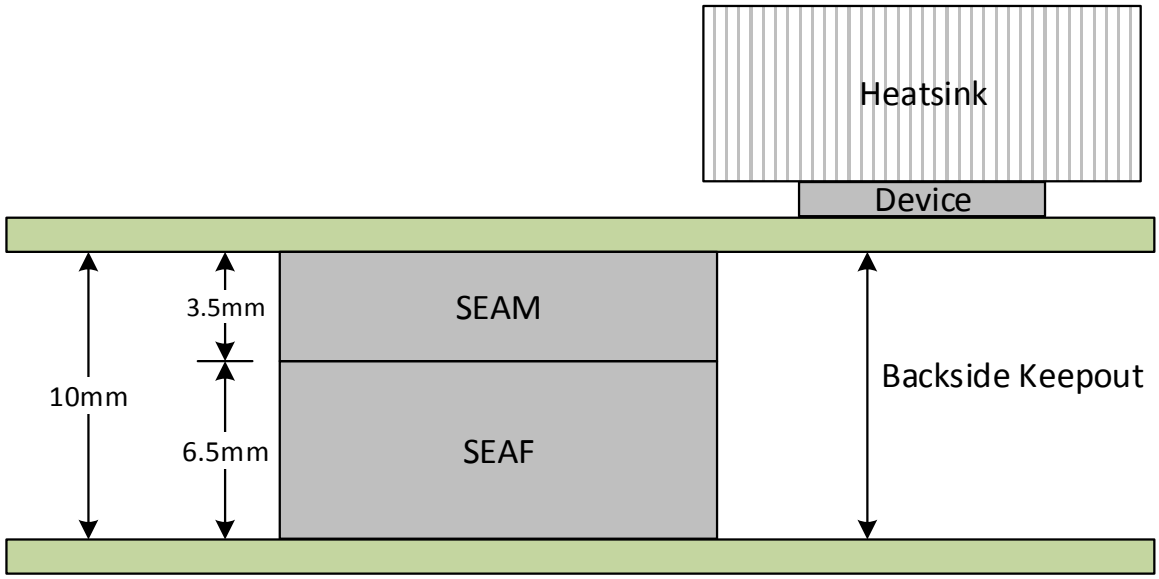

Figure 2: Connector stackup example

2.1 Signal Definitions

Table 2 defines the signals used in the NIC mezzanine interface.

Table 2: NIC mezzanine connector signal definitions

Bus type	I/O	Logic	Description
P3E_CPU1_LAN_RX_DP/N[7:0]	O	Current-mode logic (CML)	PCIe Gen3 from the NIC mezzanine card to the CPU
P3E_CPU1_LAN_TX_DP/N[7:0]	I	CML	PCIe Gen3 from the CPU to the NIC mezzanine card
CLK_100M_NIC_PE_DP/N	I	CML	100MHz PCIe Clock
PCIE_RESET_N	I	3.3V	PCIe Reset
MEZZ_PRESENT_N	O	3.3V	Mezz Present – Should be GND on mezzanine card
NWK_1_TX[0]P/N	O	CML	Port 1 GbE Transmit from mezzanine card to motherboard
NWK_1_RX[0]P/N	I	CML	Port 1 GbE Receive from motherboard to mezzanine card
NWK_2_TX[0]P/N	O	CML	Port 2 GbE Transmit from mezzanine card to motherboard
NWK_2_RX[0]P/N	I	CML	Port 2 GbE Receive from motherboard to mezzanine card
NCSI_TXD[1:0]	O	Low Voltage Transistor Transistor Logic (LVTTL)	Network Connectivity Status Indicator (NCSI) Transmit Data[1:0] from NIC to Baseboard Management Controller (BMC)
NCSI_RXD[1:0]	I	LVTTL	NCSI Transmit Data[1:0] from BMC to NIC
NCSI_CLK_IN	I	LVTTL	NCSI Input Clock
NCSI_CRS_DV	I	LVTTL	NCSI Receive Data Valid Connects to Management Controller Transmit Enable
NCSI_RX_ER	I	LVTTL	NCSI Receive Error
NWK1_ACT_LED	O	3.3V	Port 1 Activity LED
NWK1_LINK_LED	O	3.3V	Port 1 Link LED
NWK2_ACT_LED	O	3.3V	Port 2 Activity LED
NWK2_LINK_LED	O	3.3V	Port 2 Link LED
SMB_ALERT_N	O	3.3V	I2C Alert from NIC mezzanine card to BMC

Bus type	I/O	Logic	Description
NWK1_PRESENT_N	I	3.3V	Port 1 Cable Present Indicator
NWK1_I2C_SDA	I/O	3.3V	Port1 I2C Data to Cable
NWK1_I2C_CLK	O	3.3V	Port 1 I2C Clock to Cable
NWK2_PRESENT_N	I	3.3V	Port 2 Cable Present Indicator
NWK2_I2C_SDA	I/O	3.3V	Port 2 I2C Data to Cable
NWK2_I2C_SCL	O	3.3V	Port 2 I2C Clock to Cable
PCIE_WAKE_N	O	3.3V	PCIe Wake
SMB_SCL	I	3.3V	I2C to BMC
SMB_SDA	I/O	3.3V	I2C to BMC
NIC_MEZZ_ID[1:0]	O	3.3V	NIC mezzanine ID , connected to BMC on motherboard NIC_MEZZ_ID[1:0] definition is: 0:0 = Mellanox 10G 0:1 = Custom Mellanox 10G 1:0 = Emulex 10G 1:1 = Future 10G
P3V3	I	3.3V	3.3V Input Power
P3V3_AUX	I	3.3V	3.3V Aux Input Power
P12V_AUX	I	12V	12V Input Power
Ground			Ground pins

2.2 Connector Pinout

Table 3 lists the pinout for the 160-pin NIC mezzanine connector.

Table 3: NIC mezzanine connector pinout

1	GND	PCIE_RESET_N	GND	PCIE_WAKE_N	NIC_MEZZ_ID0	SMB_ALERT_N	GND	CLK_100M_NIC_PE_DP	8
9	P3E_CPU1_LAN_TX_DP<0>	GND	P3E_CPU1_LAN_RX_DP<7>	GND	SMB_SCL	NWK2_PRESENT_N	GND	CLK_100M_NIC_PE_DN	16
17	P3E_CPU1_LAN_TX_DN<0>	GND	P3E_CPU1_LAN_RX_DN<7>	GND	SMB_SDA	GND	MEZZ_PRESENT_N	GND	24
25	GND	P3E_CPU1_LAN_TX_DP<1>	GND	P3E_CPU1_LAN_RX_DP<6>	GND	NCSI_TXD_0	GND	NWK_2_TX0P	32
33	GND	P3E_CPU1_LAN_TX_DN<1>	GND	P3E_CPU1_LAN_RX_DN<6>	GND	NCSI_TXD_1	GND	NWK_2_TX0N	40
41	P3E_CPU1_LAN_TX_DP<2>	GND	P3E_CPU1_LAN_RX_DP<5>	GND	NCSI_RXD_0	GND	NWK_2_TX1P	GND	48
49	P3E_CPU1_LAN_TX_DN<2>	GND	P3E_CPU1_LAN_RX_DN<5>	GND	NCSI_RXD_1	GND	NWK_2_TX1N	GND	56
57	GND	P3E_CPU1_LAN_TX_DP<3>	GND	P3E_CPU1_LAN_RX_DP<4>	GND	NCSI_TX_EN	GND	NWK_2_TX2P	64
65	GND	P3E_CPU1_LAN_TX_DN<3>	GND	P3E_CPU1_LAN_RX_DN<4>	GND	NCSI_CLK_IN	GND	NWK_2_TX2N	72
73	P3E_CPU1_LAN_TX_DP<4>	GND	P3E_CPU1_LAN_RX_DP<0>	GND	NCSI_CRS_DV	GND	NWK_2_TX3P	GND	80
81	P3E_CPU1_LAN_TX_DN<4>	GND	P3E_CPU1_LAN_RX_DN<0>	GND	NCSI_RX_ER	GND	NWK_2_TX3N	GND	88
89	GND	P3E_CPU1_LAN_TX_DP<5>	GND	P3E_CPU1_LAN_RX_DP<1>	GND	NWK_1_ACT_LED	GND	NWK_2_RX0P	96
97	GND	P3E_CPU1_LAN_TX_DN<5>	GND	P3E_CPU1_LAN_RX_DN<1>	GND	NWK_1_LINK_LED	GND	NWK_2_RX0N	104
105	P3E_CPU1_LAN_TX_DP<6>	GND	P3E_CPU1_LAN_RX_DP<2>	GND	NWK_2_ACT_LED	GND	NWK_2_RX1N	GND	112
113	P3E_CPU1_LAN_TX_DN<6>	GND	P3E_CPU1_LAN_RX_DN<2>	GND	NWK_2_LINK_LED	GND	NWK_2_RX1P	GND	120
121	GND	P3E_CPU1_LAN_TX_DP<7>	GND	P3E_CPU1_LAN_RX_DP<3>	GND	NWK_1_TX0P	GND	NWK_2_RX2N	128
129	NWK2_I2C_SDA	P3E_CPU1_LAN_TX_DN<7>	GND	P3E_CPU1_LAN_RX_DN<3>	GND	NWK_1_TX0N	GND	NWK_2_RX2P	136
137	NWK2_I2C_SCL	GND	NIC_MEZZ_ID1	GND	NWK_1_RX0P	GND	NWK_2_RX3N	GND	144
145	P12V_AUX	P3V3_AUX	P3V3	GND	NWK_1_RX0N	GND	NWK_2_RX3P	NWK1_I2C_SDA	152
153	P12V_AUX	P3V3_AUX	P3V3	P3V3	GND	NWK1_PRESENT_N	GND	NWK1_I2C_SCL	160

2.3 Retimers/Repeaters

The blade motherboard uses 10GbE – Texas Instruments DS110DF111 (or equivalent) single-channel retimers/repeaters for interfacing to a passive cable.

2.4 Power

Table 4 shows the mezzanine power rates for each of the power rails. Note that the maximum power consumption for of the mezzanine card is 25W.

Table 4: NIC mezzanine power ratings

Power rails	Amps/pin (at 40°C)	Total number of pins	Power load capacity by connector pins (W)	Motherboard limited power budget (W)
12V_AUX	2A	2	43.2W	25.2W
3.3V_AUX	2A	2	11.88W	1.2375W
3.3V	2A	3	17.82W	9.9W
Total power budget per mezzanine card				25W

Figure 3 shows the mezzanine power-up sequence. Note that the reset will conform to the PCIe reset specification.

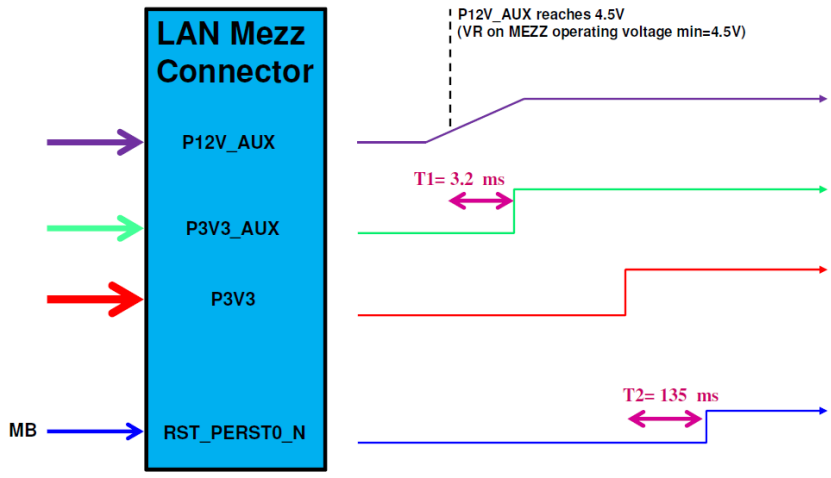

Figure 3: NIC mezzanine power-up sequence

3 Mechanical Control Outline

Figure 4 shows the mechanical outline requirements for the NIC mezzanine card.

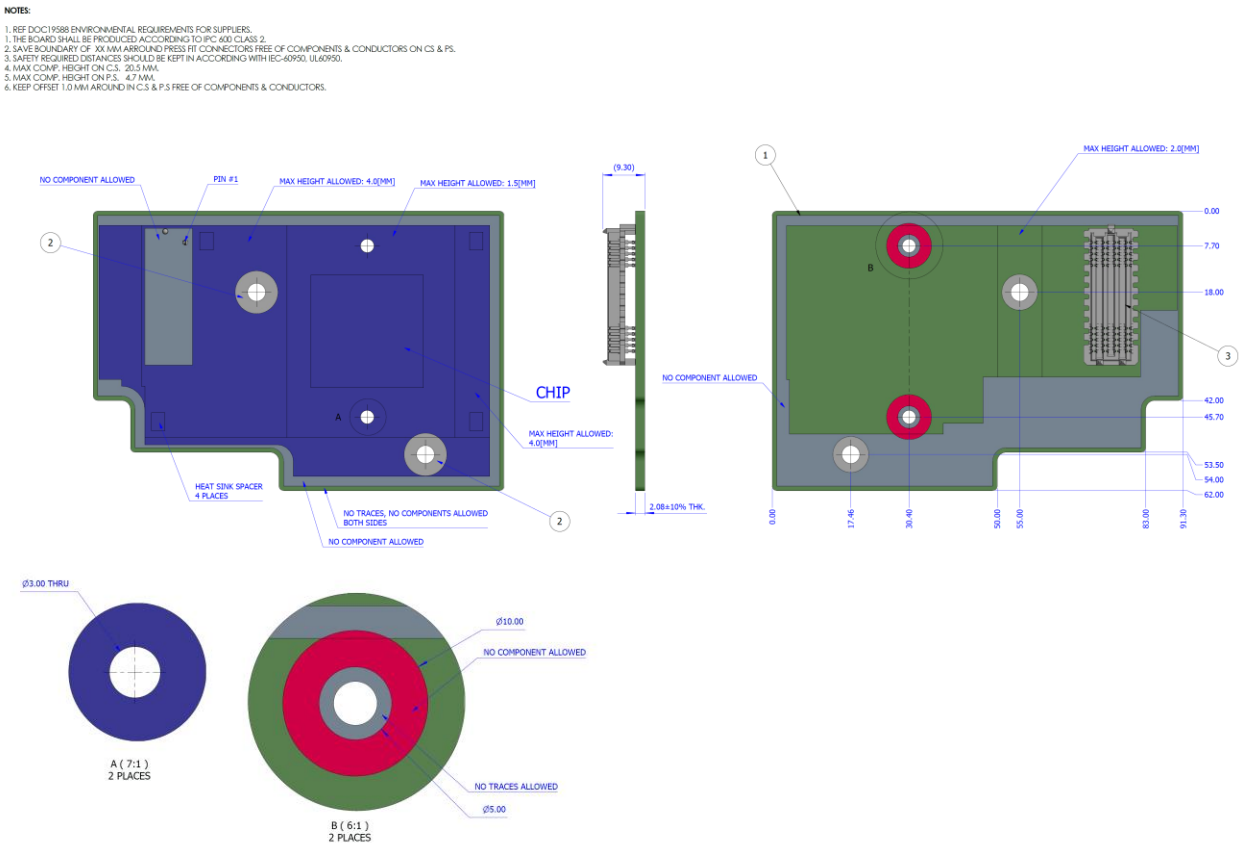

Figure 4: Mechanical control outline for the NIC mezzanine card

4 CAD Model Capture

Figure 5 shows the NIC mezzanine card installed in the blade assembly and indicates the direction of airflow.

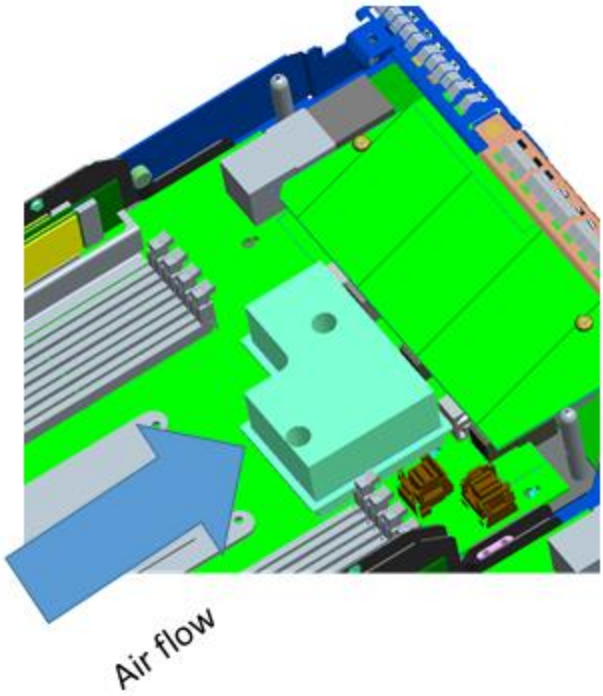

Figure 5: NIC mezzanine, installed

5 Thermal

The mezzanine card is presented with an approaching airflow temperature of 70°C and a flow rate of 0.61m/s. No downstream airflow requirements are placed on the mezzanine card.

6 Standoffs and Screws for Assembly

The standoffs that hold the NIC mezzanine card on the motherboard are soldered, therefore the NIC mezzanine card manufacturer needs to supply only two 6-32 screws with the mezzanine cards. The screws are to be provided loose in a bag and secured during assembly of the card to the motherboard.