OPEN

Compute Project

Open CloudServer
SAS mezzanine I/O specification
V1.0

Authors:

Mark Shaw, Director of Hardware Engineering, Microsoft

Martin Goldstein, Principal Systems Architect, Microsoft

1 Revision History

Date	Name	Description
1/28/2014		Version 1.0

2 Scope

This document focuses on the Open CloudServer Storage Attached SCSI (SAS) mezzanine card.

3 Contents

4 Overview

This document outlines specifications for the Open CloudServer Storage Attached SCSI (SAS) mezzanine card. The SAS card interfaces to the processor via PCI-Express (PCIe) Gen3 x8 I/O bus and to the motherboard with 8 channels 6 Gigabit SAS.

Figure 1 shows a block diagram of the connectivity of the SAS mezzanine I/O.

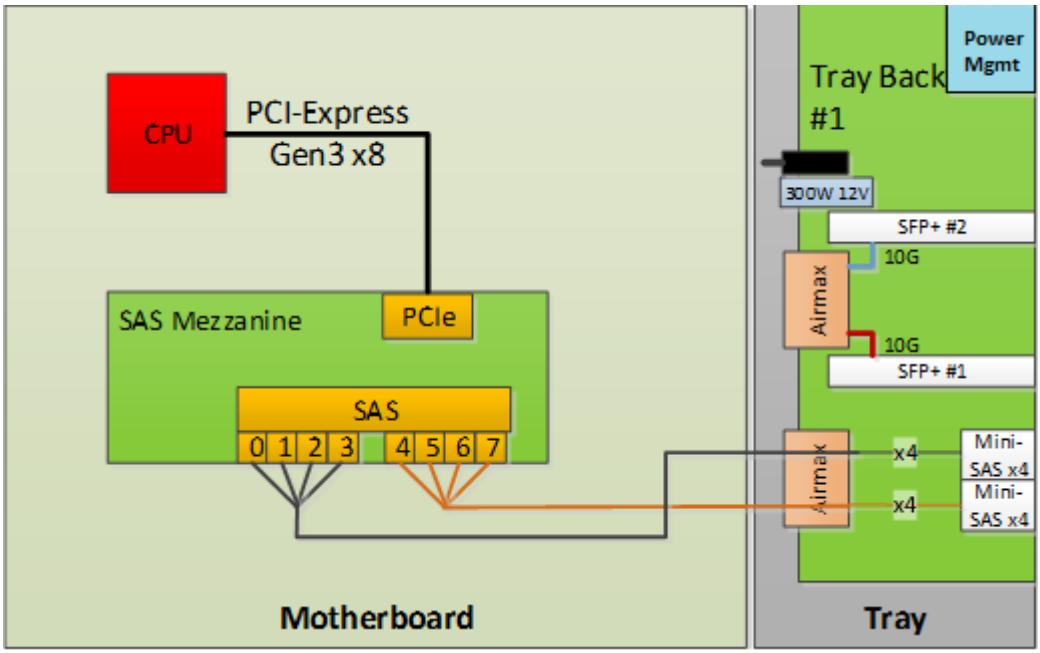

Figure 1. SAS Mezzanine I/O connectivity

5 Signaling Interface

The SAS mezzanine interface has been defined to provide high bandwidth and a flexible interface. The card receives a PCI-Express Gen3 x8 bus from the CPU.

5.1 Connectors

Table 1 shows the connector manufacturer part numbers (MPNs) for the SAS mezzanine cards.

Table 1. Connector Part Numbers, SAS Mezzanine Card

Manufacturer	Card connector MPN	Motherboard connector MPN
Samtec	SEAM-20-03.5-S-08-2-A-K-TR	SEAF-20-06.5-S-08-2-A-K-TR

Manufacturer	Card connector MPN	Motherboard connector MPN
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 are placed on the top side of the printed circuit board (PCB), as shown in Figure 2.

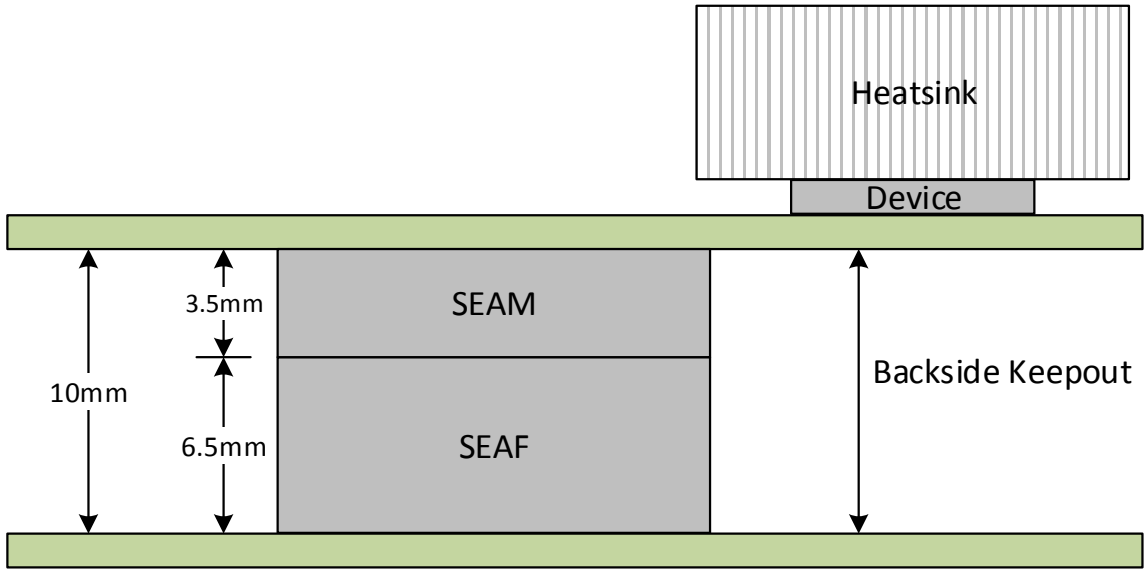

Figure 2. Connector stackup

5.2 Signal Definitions

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

Table 2. NIC Mezzanine Connector Signal Definitions

Bus type	IO	Logic	Definition
P3E_CPU1_SAS_RX_DP/N[7:0]	O	CML	PCIe Gen3 from the SAS Mezz to the CPU
P3E_CPU1_SAS_TX_DP/N[7:0]	I	CML	PCIe Gen3 from the CPU to the SAS Mezz
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
SAS_TD_CH[7:0]P/N	O	CML	SAS transmit from Mezz to motherboard

Bus type	IO	Logic	Definition
SAS_RD_CH[3:0]P/N	I	CML	SAS receive from motherboard to Mezz
SMB_ALERT_N	O	3.3V	I2C alert from SAS Mezz to baseboard management controller (BMC)
SMB_SCL	I	3.3V	I2C to BMC
SMB_SDA	I/O	3.3V	I2C to BMC
SAS_MEZZ_ID[2:0]	O	3.3V	NIC Mezz ID – connected to BMC on motherboard
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

5.3 Connector Pinout

Table 3 shows the SAS mezzanine connector interfaces and pinout numbers.

Table 3. SAS Mezzanine Connector Interface

1	MEZZ_PRSNT_N	GND	IRQ_LVC3_WAKE_N	GND	RST_PERST0_N	GND	FM_SMB_ALERT_N	GND	8
9	GND	P3e_Cpu1_Mezz_Tx_Dp<7>	GND	P3e_Cpu1_MEZZ_Rx_Dp<7>	GND	SAS_TD_CH1N	GND	SAS_RD_CH1N	16
17	GND	P3e_Cpu1_Mezz_Tx_Dn<7>	GND	P3e_Cpu1_MEZZ_Rx_Dn<7>	GND	SAS_TD_CH1P	GND	SAS_RD_CH1P	24
25	P3e_Cpu1_Mezz_Tx_Dp<6>	GND	P3e_Cpu1_MEZZ_Rx_Dp<6>	GND	SAS_TD_CH2N	GND	SAS_RD_CH2N	GND	32
33	P3e_Cpu1_Mezz_Tx_Dn<6>	GND	P3e_Cpu1_MEZZ_Rx_Dn<6>	GND	SAS_TD_CH2P	GND	SAS_RD_CH2P	GND	40
41	GND	P3e_Cpu1_Mezz_Tx_Dp<5>	GND	P3e_Cpu1_MEZZ_Rx_Dp<5>	GND	SAS_TD_CH3N	GND	SAS_RD_CH3N	48
49	GND	P3e_Cpu1_Mezz_Tx_Dn<5>	GND	P3e_Cpu1_MEZZ_Rx_Dn<5>	GND	SAS_TD_CH3P	GND	SAS_RD_CH3P	56

57	P3e_Cpu1_Mezz_Tx_Dp<4>	GND	P3e_Cpu1_MEZZ_Rx_Dp<4>	GND	SAS_TD_CH4N	GND	SAS_RD_CH4N	GND	64
65	P3e_Cpu1_Mezz_Tx_Dn<4>	GND	P3e_Cpu1_MEZZ_Rx_Dn<4>	GND	SAS_TD_CH4P	GND	SAS_RD_CH4P	GND	72
73	GND	P3e_Cpu1_Mezz_Tx_Dp<3>	GND	P3e_Cpu1_MEZZ_Rx_Dp<3>	GND	SAS_TD_CH5N	GND	SAS_RD_CH5N	80
81	GND	P3e_Cpu1_Mezz_Tx_Dn<3>	GND	P3e_Cpu1_MEZZ_Rx_Dn<3>	GND	SAS_TD_CH5P	GND	SAS_RD_CH5P	88
89	P3e_Cpu1_Mezz_Tx_Dp<2>	GND	P3e_Cpu1_MEZZ_Rx_Dp<2>	GND	SAS_TD_CH6N	GND	SAS_RD_CH6N	GND	96
97	P3e_Cpu1_Mezz_Tx_Dn<2>	GND	P3e_Cpu1_MEZZ_Rx_Dn<2>	GND	SAS_TD_CH6P	GND	SAS_RD_CH6P	GND	104
105	GND	P3e_Cpu1_Mezz_Tx_Dp<1>	GND	P3e_Cpu1_MEZZ_Rx_Dp<1>	GND	SAS_TD_CH7N	GND	SAS_RD_CH7N	112
113	GND	P3e_Cpu1_Mezz_Tx_Dn<1>	GND	P3e_Cpu1_MEZZ_Rx_Dn<1>	GND	SAS_TD_CH7P	GND	SAS_RD_CH7P	120
121	P3e_Cpu1_Mezz_Tx_Dp<0>	GND	P3e_Cpu1_MEZZ_Rx_Dp<0>	GND	SAS_TD_CH8N	GND	SAS_RD_CH8N	GND	128
129	P3e_Cpu1_Mezz_Tx_Dn<0>	GND	P3e_Cpu1_MEZZ_Rx_Dn<0>	GND	SAS_TD_CH8P	GND	SAS_RD_CH8P	GND	136
137	GND	Spare	GND	SAS_MEZZ_ID0	GND	CLK_100M_MEZZ_DP	GND	SMB_MEZZ_3V3STBY_CLK	144
145	P12V_MEZZ	P3V3_AUX	P3V3	Spare	GND	CLK_100M_MEZZ_DN	GND	SMB_MEZZ_3V3STBY_DATA	152
153	P12V_MEZZ	P3V3_AUX	P3V3	P3V3	SAS_MEZZ_ID1	GND	SAS_MEZZ_ID2	GND	160

6 Power

Table 4 shows the SAS mezzanine power ratings for the rail.

Table 4. LAN Mezzanine Power Ratings

Power rails	Amps/pin (at 40°C)	Total number of pins	Budget (W)
12V_Mezz	2	2	43.2W
3.3V_AUX	2	2	11.88W
3.3V	2	3	17.82W
Total power budget per mezzanine card)			25W

Figure 3. SAS mezzanine power-up sequence shows the SAS mezzanine power-up sequence.

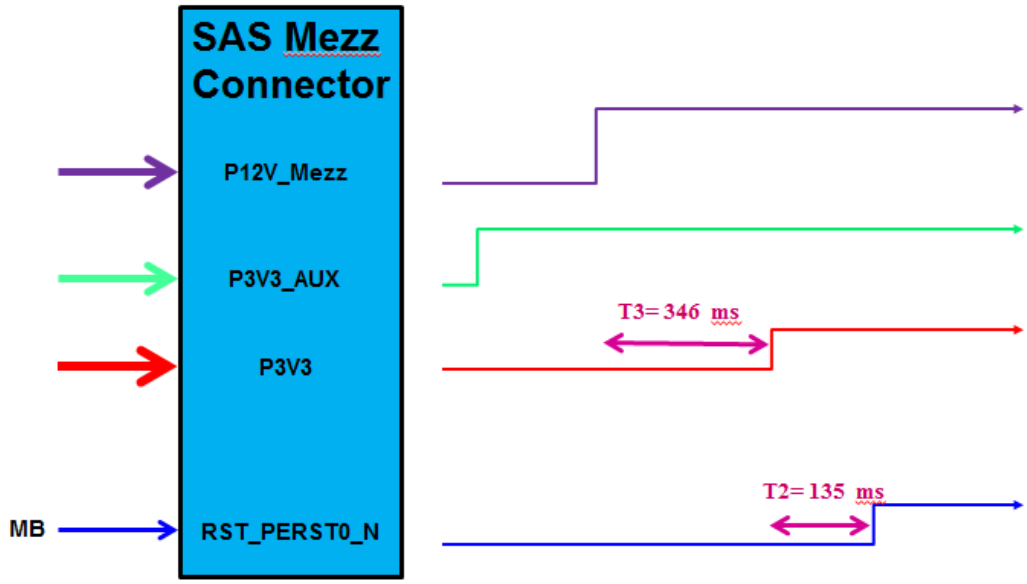

Figure 3. SAS mezzanine power-up sequence

Note that the maximum power consumption of the mezzanine cards is 25W. The reset signal conforms to the PCI Express reset specifications.

7 Mechanical

The storage controller mezzanine interface has been defined to enable an 8-channel, 6G or a 12G SAS I/O controller. The card receives a PCI-Express x8 bus from the chipset. Note that the current version of the tray only supports 6G signaling rates.

Figure 4 shows the SAS mezzanine dimensions. Note that PCB thickness can vary as

long as the keep-in volumes are not violated.

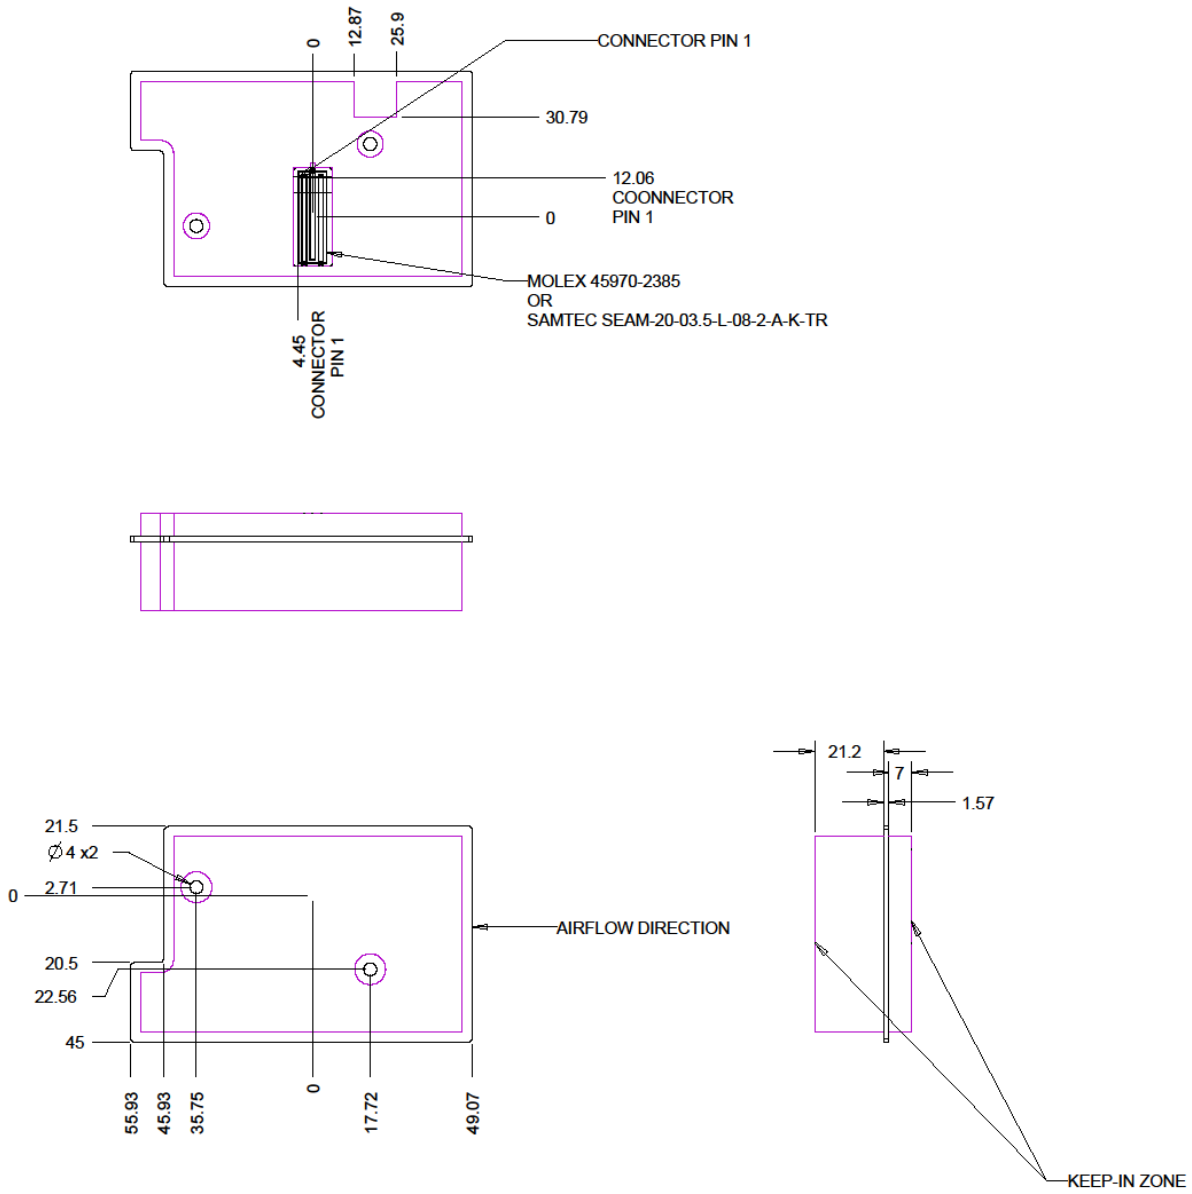

Figure 4. SAS mezzanine dimensions

8 Thermal

The SAS mezzanine cards are designed to be located at the extreme downstream position of the server, so the air is heated by all of the upstream components before reaching the cards. The direction of air flow is shown in Figure 4.

Table 5 shows the worst-case environmental conditions that can be expected at the SAS mezzanine card inlet. The thermal solution and component selection should be sufficient for these conditions.

Table 5. Environmental Operating Conditions, SAS Mezzanine Card

Variable	Worst case operating condition
Approaching airflow rate	200 ft/min, uniform
Maximum allowable pressure drop across card	0.040 "H2O max at 200 ft/min
Approaching airflow temperature	68^{O}C, uniform

9 Appendix: Commonly Used Acronyms

This section provides definitions of acronyms used in the system specifications.

ACPI – advanced configuration and power interface

AHCI – advanced host controller interface

AHJ – authority having jurisdiction

ANSI – American National Standards Institute

API – application programming interface

ASHRAE – American Society of Heating, Refrigerating and Air Conditioning Engineers

ASIC – application-specific integrated circuit

BCD – binary-coded decimal

BIOS – basic input/output system

BMC – baseboard management controller

CFM – cubic feet per minute (measure of volume flow rate)

CM – Chassis Manager

CMOS – complementary metal–oxide–semiconductor

COLO – co-location

CTS – clear to send

DCMI – Data Center Manageability Interface

DDR3 – double data rate type 3

DHCP – dynamic host configuration protocol

DIMM – dual inline memory module

DPC – DIMMs per memory channel

DRAM – dynamic random access memory

DSR – data set ready

DTR – data terminal ready

ECC – error-correcting code

EEPROM - electrically erasable programmable read-only memory

EIA – Electronic Industries Alliance

EMC – electromagnetic compatibility

EMI – electromagnetic interference

FRU – field replaceable unit

FTP – file transfer protocol

GPIO – general purpose input output

GUID – globally unique identifier

HBI – high business intelligence

HCK – Windows Hardware Certification Kit

HMD – hardware monitoring device

HT – hyperthreading

I²C – inter-integrated circuit

IBC – international building code

IDE – integrated development environment

IEC – International Electrotechnical Commission

IOC – I/O controller

IPMI – intelligent platform management interface

IPsec – IP security

ITPAC – IT pre-assembled components

JBOD – "just a bunch of disks"

KCS – keyboard controller style

L2 – layer 2

LAN – local area network

LFF – large form factor

LPC – low pin count

LS – least significant

LUN – logical unit number

MAC – media access control

MDC – modular data center containers

MLC – multi-level call

MTBF – mean time between failures

MUX - multiplexer

NIC – network interface card

NUMA – non-uniform memory access

OOB – out of band

OSHA - Occupational Safety & Health Administration

OTS – off the shelf

PCB – printed circuit board

PCIe – peripheral component interconnect express

PCH – platform control hub

PDB – power distribution backplane

PDU – power distribution unit

Ph-ph – phase to phase

Ph-N – phase to neutral

PNP – plug and play

POST – power-on self-test

PSU – power supply unit

PWM – pulse-width modulation

PXE – preboot execution environment

QDR – quad data rate

QFN – quad flat package no-lead

QPI – Intel QuickPath Interconnect

QSFP – Quad small form-factor pluggable

RAID – redundant array of independent disks

REST - representational state transfer

RM – Rack Manager

RMA – remote management agent

ROC – RAID-on-chip controller

RSS – receive-side scaling

RTS – request to send

RU – rack unit

RxD – received data

SAS – serial-attached small computer system interface (SCSI)

SATA – serial AT attachment

SCK – serial clock

SCSI – small computer system interface

SDA – serial data signal

SDR – sensor data record

SFF – small form factor

SFP - small form-factor pluggable

SMBUS – systems management bus

SMBIOS – systems management BIOS

SOL – serial over LAN

SPI – serial peripheral interface

SSD – solid-state drive

TB – tray backplane

TDP – thermal design power

TB – tray backplane

TOR – top of rack

TPM – trusted platform module

TxD – transmit data

U – rack unit

UART – universal asynchronous receiver/transmitter

UEFI – unified extensible firmware interface

UL – Underwriters Laboratories

UPS – uninterrupted power supply

Vpp – voltage peak to peak

WMI – Windows Management Interface