SSI Compute Blade Mezzanine Specification

November 2010

Revision 1.0.1

Important Information and Disclaimers:

1. THE SERVER SYSTEM INFRASTRUCTURE PROMOTERS (“SSI PROMOTERS”) MAKE NO WARRANTIES WITH REGARD TOTHIS SSI SPECIFICATION (“SPECIFICATION”), AND IN PARTICULAR DOES NOT WARRANT OR REPRESENT THAT THIS SPECIFICATION OR ANY PRODUCTS MADE IN CONFORMANCE WITH IT WILL WORK IN THE INTENDED MANNER. NOR WILL SSI PROMOTERS ASSUME ANY RESPONSIBILITY FOR ANY ERRORS THAT THE SPECIFICATION MAY CONTAIN OR HAVE ANY LIABILITIES OR OBLIGATIONS FOR DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, INDIRECT, PUNITIVE, OR CONSEQUENTIAL DAMAGES WHETHER ARISING FROM OR IN CONNECTION WITH THE USE OF THIS SPECIFICATION IN ANY WAY.

2. NO REPRESENTATIONS OR WARRANTIES ARE MADE THAT ANY PRODUCT BASED INWHOLE OR PART ON THE ABOVE SPECIFICATION WILL BE FREE FROM DEFECTS OR SAFE FOR USE FOR ITS INTENDED PURPOSE. ANY PERSON MAKING, USING OR SELLING SUCH PRODUCT DOES SO AT HIS OR HER OWN RISK.

3. THE USER OF THIS SPECIFICATION HEREBY EXPRESSLY ACKNOWLEDGES THAT THE SPECIFICATION IS PROVIDED AS IS, AND THAT THE SSI PROMOTERS MAKE NO REPRESENTATIONS, EXTENDS NO WARRANTIES OF ANY KIND EITHER EXPRESS OR IMPLIED ORAL OR WRITTEN, INCLUDING ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTY OR REPRESENTATION THAT THE SPECIFICATION OR ANY PRODUCT OR TECHNOLOGY UTILIZING THE SPECIFICATION OR ANY SUBSET OF THE SPECIFICATION WILL BE FREE FROM ANY CLAIMS OF INFRINGEMENT OF INTELLECTUAL PROPERTY, INCLUDING PATENTS, COPYRIGHTS AND TRADE SECRETS NOR DO THE SSI PROMOTERS ASSUME ANY OTHER RESPONSIBILITIES WHATSOEVER WITH RESPECT TO THE SPECIFICATION OR SUCH PRODUCTS.

4. A NON-EXCLUSIVE COPYRIGHT LICENSE IS HEREBY GRANTED TO REPRODUCE THIS SPECIFICATION FOR ANY PURPOSE PROVIDED THIS “IMPORTANT INFORMATION AND DISCLAIMERS SECTION (PARAGRAPHS 1-6) IS PROVIDED IN WHOLE.

5. UPON REQUEST FROM AN ADOPTER, THE SSI PROMOTERS WILL GRANT A NON-EXCLUSIVE, WORLD-WIDE LICENSE UNDER ANY NECESSARY CLAIMS, DEFINED IN THE ADOPTERS AGREEMENT, TO MAKE, HAVE MADE, USE, IMPORT, SELL, OFFER TO SELL, AND OTHERWISE DISTRIBUTE AND DISPOSE OF COMPLIANT PORTIONS, DEFINED IN THE ADOPTERS AGREEMENT, PROVIDED THAT SUCH LICENSE NEED NOT EXTEND TO ANY PART OR FUNCTION OF A PRODUCT IN WHICH A COMPLIANT PORTION IS INCORPORATED THAT IS NOT ITSELF PART OF THE COMPLIANT PORTION. SUCH LICENSE WILL BE GRANTED ON REASONABLE AND NONDISCRIMINATORY (“RAND”) TERMS AND MAY BE CONDITIONED UPON ADOPTER’S GRANT OF A RECIPROCAL LICENSE TO THE SSI PROMOTERS.

6. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY OTHER INTELLECTUAL PROPERTY RIGHTS IS GRANTED.

Revision 1.0.1

*Product names are trademarks, registered trademarks, or servicemarks of their respective owners.

Contents

Figures

Tables

Revision History

The following table lists the revision schedule based on revision number and development stage of the product.

Revision	Project Document State	Date
1.0.0	Initial public release.	September 16, 2009
1.0.1	Bug updates: 302, 382	November 4, 2010

Note: Not all revisions may be published.

This page intentionally left blank.

1 Mezzanine Architecture

1.1 Introduction

For Compute Blades that need an additional high-speed I/O fabric connection to the Flexi mid-plane interconnect, an optional mezzanine PBA may be added. The mezzanine PBA is architecturally equivalent to a standard PCIe* adapter card with the major differences being in form factor and connector interface. The software interface is unchanged.

The mezzanine PBA interface to the Compute Blade can support x16, x8, x4, and x1 PCIe connections. The interface is through two connectors: a primary mezzanine/blade connector and an optional expansion mezzanine/blade connector.

The primary interface is a x8 PCI Express* connection using a 120-pin Tyco* Flex height (or equivalent) connector. The optional expansion interface is an 80-pin connector (Tyco Flex height or equivalent) and defined for an additional x8 PCIe interface. If both connectors are implemented, the two x8 PCIe interfaces can be aggregated into a single x16 PCIe interface. The PCIe interfaces are compatible with PCI Express revision 2.

Figure 1-1 shows how the mezzanine PBA connects to the Compute Blade and to the Flexi mid-plane interconnect.

Figure 1-1: Mezzanine Connection Diagram

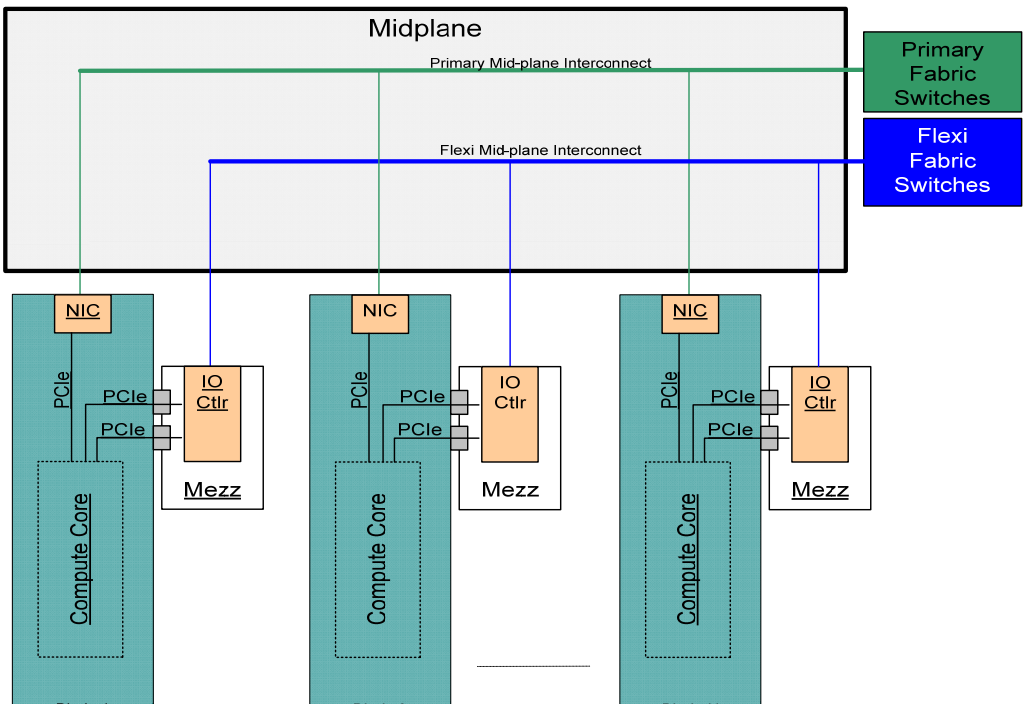

The mezzanine high-speed I/O connects directly to the Flexi mid-plane through an Airmax* VS 96-pin, right angle connector.

Two pairs of Flexi mid-plane interconnects are defined: a pair of x4 fabric links (4 channels) and a pair of x1 links (1 channel). It is required that both interconnects within a pair be implemented for redundancy. Both interconnects within a pair **shall** be of the same technology. It is not required that the technology be the same on all four interconnects, only within a pair.

It is optional that a mid-plane route a pair of links to redundant switches. A mid-plane may choose to aggregate the two links and route them to the same switch.

10G Ethernet*, Fibre Channel*, or InfiniBand* technologies can be used for these connections. Others may also be used.

1.1.1 Reference Documents

- IPMI – Intelligent Platform Management Interface Specification, v2.0 rev 1.0E3, February 16, 2006, Copyright © 2004, 2005, 2006 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.
- IPMI – Platform Management FRU Information Storage Definition, V1.0, Document revision 1.1, September 27, 1999 Copyright © 1998, 1999 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.
- IPMI – Intelligent Platform Management Bus Communications Protocol Specification, V1.0, Document revision 1.0, November 15, 1999 Copyright © 1998, 1999 Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Inc., All rights reserved.
- SSI Compute Blade Specification Rev. 1.0.0
- PCI Express Base Specification Rev. 2.0

1.1.2 Terms and Abbreviations

Table 1-1 lists terms and acronyms used in specific ways throughout this specification.

Table 1-1: Terms and Abbreviations

Term	Definition
ASHRAE	American Society of Heating, Refrigerating, and Air Conditioning Engineers.
Base Management Interface (BMI)	This is the IPMB-based management interface used by the Chassis Manager to communicate with the blade management controllers.
blade	This is a resource module that plugs into the blade chassis. A blade can provide many different types of resources to the chassis, including compute functions, storage capabilities, additional I/O interfaces and switching capabilities, and special purpose capabilities. A blade can be a single-wide module (assumed) or a double-wide module, occupying two adjacent slots in the chassis.
blade server	A system comprising a chassis, chassis resources (power, cooling, Chassis Manager), compute blades, and communication (switch) blades. The chassis may contain additional modules, such as storage.
bottom	When used in reference to a board, the end that would be on the bottom in a vertically oriented chassis.
CFM	Cubic Feet per Minute. A measure of volumetric airflow. One CFM is equivalent to 472 cubic centimeters per second.
chassis	The mechanical enclosure that consists of the mid-plane, front boards, cooling devices, power supplies, etc. The chassis provides the interface to boards, and it consists of the guide rails, alignment, handle interface, face plate mounting hardware, and mid-plane interface.
chassis ground	A safety ground and earth return that is connected to the chassis metal and available to all PBAs.
Chassis Management Module (CMM)	Dedicated intelligent chassis module that hosts the Chassis Manager functionality.
Chassis Manager (CM)	Set of logical functions for hardware management of the chassis. This may be implemented by one or more dedicated Chassis Management Modules or by one or more blade management controllers and/or payload processors.
cold start	Cold start is the time when blades receive the payload power for the first time.
component side 1	When used in reference to a PBA, the side on which the tallest electronic components would be mounted.
component side 2	When used in reference to a PBA, the side normally reserved for making solder connections with through-hole components on Component Side 1, but on which low-height electronic components may also be mounted.
creepage	Surface distance required between two electrical components.
face plate	The front-most element of a PBA, perpendicular to the PBA, that serves to mount connectors, indicators, controls, and mezzanines.
guide rail	Provides for the front board guidance feature in a slot.

Term	Definition
handle	An item or part used to insert or extract blades into and out of chassis.
Intelligent Platform Management Bus (IPMB)	IPMB is an I2C-based bus that provides a standardized interconnection between managed modules within a chassis. ftp://download.intel.com/design/servers/ipmi/ipmb1010ltd.pdf
Intelligent Platform Management Interface (IPMI)	IPMI v2.0 R1.0 specification defines a standardized, abstracted interface to the platform management subsystem of a computer system. ftp://download.intel.com/design/servers/ipmi/IPMIv2_0rev1_0.pdf
interconnect channel	An interconnect channel comprises two pairs of differential signals. One pair of differential signals for transmit and another pair of differential signals for receive.
LFM	Linear Feet per Minute. A measure of air velocity. One LFM is equivalent to 0.508 centimeters per second.
logic ground	Chassis-wide electrical net used on blades and mid-planes as a reference and return path for logic-level signals that are carried between boards.
managed module	Any component of the system that is addressable for management purposes via the specified management interconnect and protocol. A managed module is interfaced directly to the chassis BMI.
Management Controller	This is an intelligent, embedded microcontroller that provides management functionality for a blade or other chassis module.
may	Indicates flexibility of choice with no implied preference.
MBSTC	Generic term Midplane Blade-Switch test card for measurement of electrical parameters. MBSTC is super set of MBTC and MSTC.
MBTC-xx	Blade test card replacing a switch for measurement of electrical parameters with accompanying index xx
mezzanine	The mezzanine is a PBA that installs on a blade PBA horizontally. It provides additional functionality on the blade PBA and provides electrical interface between the blade PBA and the mid-plane PBA. Both the blade PBA and mezzanine PBA are contained inside the blade module.
mid-plane	Equivalent to a system backplane. This is a PBA that provides the common electrical interface for each blade in the chassis and on both the front and back of the PBA.
module	A physically separate chassis component which may be independently replaceable (e.g., a blade or cooling module) or attached to some other component (e.g., a mezzanine board).
MSTC-xx	Midplane Switch test card replacing a blade for measurement of electrical parameters with accompanying index xx
open blade	A blade that conforms to the requirements defined by the Open Blade standard set of specifications.
out-of-band (OOB)	Communication between blades that does not need the host or payload to be powered on.
payload	The hardware on a blade that implements the main mission function of the blade. On a compute blade, this includes the main processors, memory, and I/O interfaces. The payload is powered separately from the blade management subsystem. Payload power is controlled by the blade management controller.
PBA	Printed board assembly. A printed circuit board that has all electronic components attached to it.

Term	Definition
PCB	Printed circuit board without components attached.
peak power	The maximum power a blade can draw for a very short period of time during a hot insertion, hot removal, or a cold start.
pitch line	Horizontal pitch line between slots.
shall	Indicates a mandatory requirement. Designers must implement such mandatory requirements to ensure interchangeability and to claim conformance with this specification. The use of **shall not** (in bold) indicates an action or implementation that is prohibited.
should	Indicates flexibility of choice with a strongly preferred implementation. The use of **should not** (in bold) indicates flexibility of choice with a strong preference that the choice or implementation be avoided.
slot	A slot defines the position of one blade in a chassis.
top	When used in reference to a blade, the end which would be on top in a vertically oriented chassis.
U	Unit of vertical height defined in IEC 60297-1 rack, shelf, and chassis height increments. 1U=44.45 mm.
WDT	Watchdog timer.

1.2 Mezzanine PBA Mechanical Specifications

The maximum mezzanine PBA dimension **shall** be 85 mm x 155 mm, as shown in Figure 1-2.

The *primary* side of the mezzanine PBA faces *away* from the Compute Blade. The *secondary* side of the mezzanine PBA faces *toward* the Compute Blade. The two Compute Blade interface connectors (1 and 2 in Figure 1-2) and the Flexi mid-plane connector (3 in Figure 1-2) attach to the secondary side of the mezzanine PBA.

Figure 1-3 illustrates a physical view of the mezzanine PBA.

Figure 1-2: Mezzanine PBA Mechanical Specification

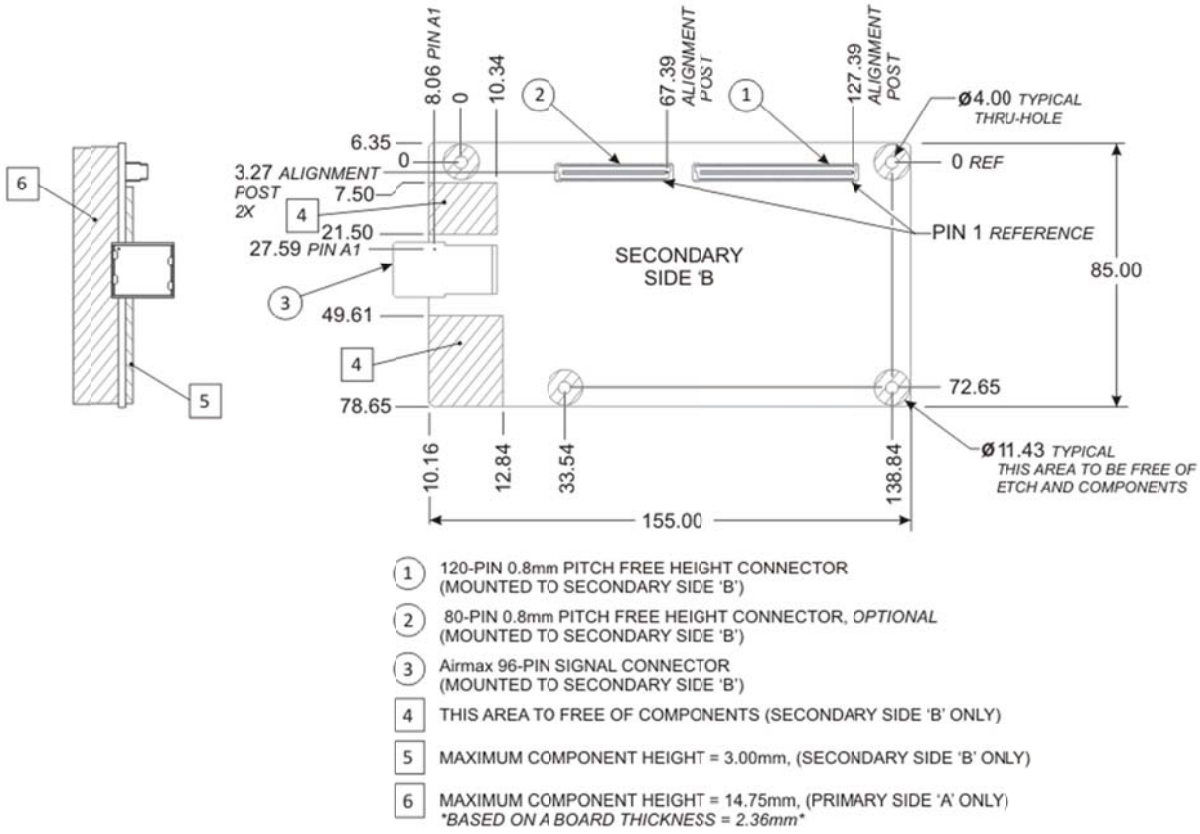

Figure 1-3: Isometric View of Mezzanine PBA

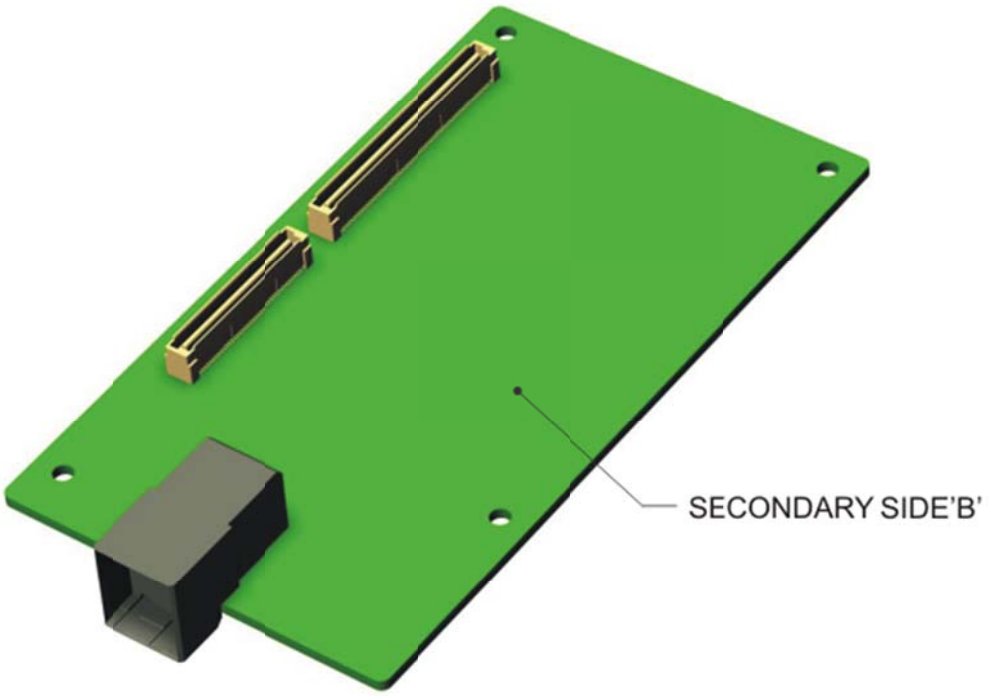

When the mezzanine PBA is plugged into the Compute Blade, there is an opportunity for interference with components on the Compute Blade PCA. Therefore, the **maximum** height for components on the secondary side of the mezzanine PBA **shall** be 3.00 mm. The area on the Compute Blade beneath the mezzanine PBA **should** be free of tall components that may interfere with components on the mezzanine PBA or obstruct airflow into the mezzanine PBA.

With a PCB thickness of 0.093" (2.36mm) [382], the maximum height for components on the primary side of the mezzanine PBA is 14.75 mm.

The mezzanine PBA to Flexi mid-plane connector (labeled as 3 in Figure 1-2) is a 96-pin Airmax VS connector, or mechanical equivalent. It is located on the secondary side of the mezzanine PBA.

The primary blade/mezzanine connector (labeled as 1 in Figure 1-2) is a 120-pin Tyco* 0.8 mm pitch Free Height connector, or mechanical equivalent. The plug is located on the secondary side of the mezzanine PBA, while the receptacle is located on the Compute Blade. The mated height is 16 mm.

The expansion blade/mezzanine connector (labeled as 2 in Figure 1-2) is an 80-pin Tyco 0.8 mm pitch Free Height connector or mechanical equivalent. The plug is located on the secondary side of the mezzanine PBA, while the receptacle is located on the Compute Blade. The mated height is 16 mm.

Minimum board thickness is 0.062".

Table 1-2 lists the part numbers for the mezzanine PBA connectors. Mechanical equivalents may be used.

Table 1-2: Connector Part Numbers

Connector	Description	Manufacturer	Part Number
96-pin Flexi-Fabric	4-pair vert, 2mm, 96-pin plug	FCI*	10084604-111LF or equivalent.
Primary Blade connector (Mezzanine Half)	120-pin dual-row plug	Tyco*	5179031-5 or equivalent.
Primary Blade Connector (Compute Blade Half)	120-pin dual-row receptacle	Tyco*	5-5179010-5 or equivalent.
Expansion Blade Connector (Mezzanine Half)	80-pin dual-row plug	Tyco*	5179031-3 or equivalent.
Expansion Blade Connector (Compute Blade Half)	80-pin dual-row receptacle	Tyco*	5-5179010-3 or equivalent.

Note: the primary and expansion blade connector part numbers use 8 micron thick gold plating. Environmental requirements of specific implementations may require thicker plating.

1.2.1 Primary Side Component Height

The *primary* side of the mezzanine PBA faces *away* from the Compute Blade.

The maximum component height must not exceed 14.75 mm with a PCB thickness of 0.093".

1.2.2 Secondary Side Component Height

The *secondary* side of the mezzanine PBA faces *toward* the Compute Blade.

When the mezzanine PBA is plugged into the Compute Blade, there is an opportunity for interference with components on the Compute Blade PCB. Therefore, the **maximum** height for components on the secondary side of the mezzanine PBA **shall** be 3.00 mm.

1.2.3 Mezzanine to Compute Blade Connector Pinout

The mezzanine PBA interfaces to the Compute Blade through a primary blade/mezzanine connector and an expansion blade/mezzanine connector.

Table 1-3 and Table 1-4 show the primary blade/mezzanine connector pinout and signal definitions. The connector contains a x8 PCIe interface, plus all power and side band signals. Table 1-5 and Table 1-5 show the expansion blade/mezzanine connector pinout and signal definitions. This connector has an additional x8 PCIe interface.

The direction of the interface is with respect to the Compute Blade. For example, RX goes to the receiver on the Compute Blade.

Table 1-3: Primary Blade / Mezzanine Connector

P5V	1		2	P5V
GND	3		4	GND
P3V3	5		6	P3V3
P3V3	7		8	P3V3
P3V3	9		10	P3V3
GND	11		12	GND
P3V3AUX	13		14	P3V3AUX
P3V3AUX	15		16	P3V3AUX
SMB_SDA	17		18	SMB_SCL
HSC0_LNK_LED_N	19		20	HSC0_ACT_LED_N
HSC1_LNK_LED_N	21		22	HSC1_ACT_LED_N
HSC2_LNK_LED_N	23		24	HSC2_ACT_LED_N
HSC3_LNK_LED_N	25		26	HSC3_ACT_LED_N
GND	27		28	WAKE_N
RSVD	29		30	GND
RSVD	31		32	GND
GND	33		34	PCIe_0_A_TXP
GND	35		36	PCIe_0_A_TXN
PCIe_0_A_RXP	37		38	GND
PCIe_0_A_RXN	39		40	GND
GND	41		42	PCIe_0_B_TXP
GND	43		44	PCIe_0_B_TXN
PCIe_0_B_RXP	45		46	GND
PCIe_0_B_RXN	47		48	GND
GND	49		50	PCIe_0_C_TXP
GND	51		52	PCIe_0_C_TXN
PCIe_0_C_RXP	53		54	GND
PCIe_0_C_RXN	55		56	GND
GND	57		58	PCIe_0_D_TXP
GND	59		60	PCIe_0_D_TXN
PCIe_0_D_RXP	61		62	GND
PCIe_0_D_RXN	63		64	GND
GND	65		66	PCIe_1_A_TXP
GND	67		68	PCIe_1_A_TXN
PCIe_1_A_RXP	69		70	GND
PCIe_1_A_RXN	71		72	GND
GND	73		74	PCIe_1_B_TXP
GND	75		76	PCIe_1_B_TXN
PCIe_1_B_RXP	77		78	GND
PCIe_1_B_RXN	79		80	GND
GND	81		82	PCIe_1_C_TXP
GND	83		84	PCIe_1_C_TXN
PCIe_1_C_RXP	85		86	GND
PCIe_1_C_RXN	87		88	GND
GND	89		90	PCIe_1_D_TXP
GND	91		92	PCIe_1_D_TXN
PCIe_1_D_RXP	93		94	GND
PCIe_1_D_RXN	95		96	GND
GND	97		98	Mezz_Present
GND	99		100	Reset_N
Clk0_100M_PCIE_P	101		102	GND
Clk0_100M_PCIE_N	103		104	GND
GND	105		106	RSVD
GND	107		108	RSVD
RSVD	109		110	GND
RSVD	111		112	RSVD
RSVD	113		114	RSVD
P12V	115		116	P12V
P12V	117		118	P12V
P12V	119		120	P12V

Note: Green color indicates GND (ground) connection. Yellow are reserved (RSVD) connections.

Table 1-4: Primary Blade / Mezzanine Connector Signal Definitions

Signal Name	Signal Description	Purpose	Connector Location
PCIe_0_A_TXP	PCIe TX+ of Lane A Link 0	Host connect	34
PCIe_0_A_TXN	PCIe TX- of Lane A Link 0	Host connect	36
PCIe_0_A_RXP	PCIe RX+ of Lane A Link 0	Host connect	37
PCIe_0_A_RXN	PCIe RX- of Lane A Link 0	Host connect	39
PCIe_0_B_TXP	PCIe TX+ of Lane B Link 0	Host connect	42
PCIe_0_B_TXN	PCIe TX- of Lane B Link 0	Host connect	44
PCIe_0_B_RXP	PCIe RX+ of Lane B Link 0	Host connect	45
PCIe_0_B_RXN	PCIe RX- of Lane B Link 0	Host connect	47
PCIe_0_C_TXP	PCIe TX+ of Lane C Link 0	Host connect	50
PCIe_0_C_TXN	PCIe TX- of Lane C Link 0	Host connect	52
PCIe_0_C_RXP	PCIe RX+ of Lane C Link 0	Host connect	53
PCIe_0_C_RXN	PCIe RX- of Lane C Link 0	Host connect	55
PCIe_0_D_TXP	PCIe TX+ of Lane D Link 0	Host connect	58
PCIe_0_D_TXN	PCIe TX- of Lane D Link 0	Host connect	60
PCIe_0_D_RXP	PCIe RX+ of Lane D Link 0	Host connect	61
PCIe_0_D_RXN	PCIe RX- of Lane D Link 0	Host connect	63
PCIe_1_A_TXP	PCIe TX+ of Lane A Link 1	Host connect	66
PCIe_1_A_TXN	PCIe TX- of Lane A Link 1	Host connect	68
PCIe_1_A_RXP	PCIe RX+ of Lane A Link 1	Host connect	69
PCIe_1_A_RXN	PCIe RX- of Lane A Link 1	Host connect	71
PCIe_1_B_TXP	PCIe TX+ of Lane B Link 1	Host connect	74
PCIe_1_B_TXN	PCIe TX- of Lane B Link 1	Host connect	76
PCIe_1_B_RXP	PCIe RX+ of Lane B Link 1	Host connect	77
PCIe_1_B_RXN	PCIe RX- of Lane B Link 1	Host connect	79
PCIe_1_C_TXP	PCIe TX+ of Lane C Link 1	Host connect	82
PCIe_1_C_TXN	PCIe TX- of Lane C Link 1	Host connect	84
PCIe_1_C_RXP	PCIe RX+ of Lane C Link 1	Host connect	85
PCIe_1_C_RXN	PCIe RX- of Lane C Link 1	Host connect	87
PCIe_1_D_TXP	PCIe TX+ of Lane D Link 1	Host connect	90
PCIe_1_D_TXN	PCIe TX- of Lane D Link 1	Host connect	92
PCIe_1_D_RXP	PCIe RX+ of Lane D Link 1	Host connect	93
PCIe_1_D_RXN	PCIe RX- of Lane D Link 1	Host connect	95
Clk0_100M_PCIe_P	100MHz clk +	PCIe Clk	101
Clk0_100M_PCIe_N	100MHz clk -	PCIe Clk	103
SMB_SCL	SMBus Clock	Mngt connect	18
SMB_SDA	SMBus Data	Mngt connect	17

Signal Name	Signal Description	Purpose	Connector Location
HSC_0_LNK_LED_N	HSC 0 Link LED driver (Active Low)	LED control	19
HSC_1_LNK_LED_N	HSC 1 Link LED driver(Active Low)	LED control	21
HSC_2_LNK_LED_N	HSC 2 Link LED driver(Active Low)	LED control	23
HSC_3_LNK_LED_N	HSC 3 Link LED driver(Active Low)	LED control	25
HSC_0_ACT_LED_N	HSC 0 Activity LED driver(Active Low)	LED control	20
HSC_1_ACT_LED_N	HSC 1 Activity LED driver (Active Low)	LED control	22
HSC_2_ACT_LED_N	HSC 2 Activity LED driver (Active Low)	LED control	24
HSC_3_ACT_LED_N	HSC 3 Activity LED driver (Active Low)	LED control	26
WAKE_N	PCIe WAKE_N signal	Wake on LAN	28
Reset_N	Reset signal (Active Low)	Mezz Reset	100
Mezz_PRES_N	Mezzanine Present signal (active Low)	Present indication	98
P12V	12V Power	Power	115, 116, 117, 118, 119, 120
P3V3	3.3V Power	Power	5, 6, 7, 8, 9, 10
P5V	5V Power	Power	1, 2
P3V3AUX	Auxiliary Power	Aux Power	13, 14, 15, 16
RSVD	Reserved pins	Future use	29, 31, 106, 108, 109, 111, 112, 113, 114
GND	Ground		3, 4, 11, 12, 27, 30, 32, 33, 35, 38, 40, 41, 43, 46, 48, 49, 51, 54, 56, 57,59, 62, 64, 65, 67, 70, 72, 73, 75, 78, 80, 81, 83, 86, 88, 89, 91, 94, 96, 97, 99, 102, 104, 105, 107, 110

Table 1-5: Expansion Blade / Mezzanine Connector

RSVD	1		2	GND
RSVD	3		4	GND
GND	5		6	PCIe_2_A_TXP
GND	7		8	PCIe_2_A_TXN
PCIe_2_A_RXP	9		10	GND
PCIe_2_A_RXN	11		12	GND
GND	13		14	PCIe_2_B_TXP
GND	15		16	PCIe_2_B_TXN
PCIe_2_B_RXP	17		18	GND
PCIe_2_B_RXN	19		20	GND
GND	21		22	PCIe_2_C_TXP
GND	23		24	PCIe_2_C_TXN
PCIe_2_C_RXP	25		26	GND
PCIe_2_C_RXN	27		28	GND
GND	29		30	PCIe_2_D_TXP
GND	31		32	PCIe_2_D_TXN
PCIe_2_D_RXP	33		34	GND
PCIe_2_D_RXN	35		36	GND
GND	37		38	PCIe_3_A_TXP
GND	39		40	PCIe_3_A_TXN
PCIe_3_A_RXP	41		42	GND
PCIe_3_A_RXN	43		44	GND
GND	45		46	PCIe_3_B_TXP
GND	47		48	PCIe_3_B_TXN
PCIe_3_B_RXP	49		50	GND
PCIe_3_B_RXN	51		52	GND
GND	53		54	PCIe_3_C_TXP
GND	55		56	PCIe_3_C_TXN
PCIe_3_C_RXP	57		58	GND
PCIe_3_C_RXN	59		60	GND
GND	61		62	PCIe_3_D_TXP
GND	63		64	PCIe_3_D_TXN
PCIe_3_D_RXP	65		66	GND
PCIe_3_D_RXN	67		68	GND
GND	69		70	Clk1_100M_PCIE_P
GND	71		72	Clk1_100M_PCIE_N
RSVD	73		74	GND
RSVD	75		76	GND
RSVD	77		78	RSVD
RSVD	79		80	RSVD

Note: Green indicates GND (ground) connections. Yellow indicates reserved (RSVD) connections.

Table 1-6: Expansion Blade / Mezzanine Connector Signal Definitions

Signal Name	Signal Description	Purpose	Connector Location
PCIe_2_A_TXP	PCIe TX+ of Lane A Link 2	Host connect	6
PCIe_2_A_TXN	PCIe TX- of Lane A Link 2	Host connect	8
PCIe_2_A_RXP	PCIe RX+ of Lane A Link 2	Host connect	9
PCIe_2_A_RXN	PCIe RX- of Lane A Link 2	Host connect	11
PCIe_2_B_TXP	PCIe TX+ of Lane B Link 2	Host connect	14
PCIe_2_B_TXN	PCIe TX- of Lane B Link 2	Host connect	16
PCIe_2_B_RXP	PCIe RX+ of Lane B Link 2	Host connect	17
PCIe_2_B_RXN	PCIe RX- of Lane B Link 2	Host connect	19
PCIe_2_C_TXP	PCIe TX+ of Lane C Link 2	Host connect	22
PCIe_2_C_TXN	PCIe TX- of Lane C Link 2	Host connect	24
PCIe_2_C_RXP	PCIe RX+ of Lane C Link 2	Host connect	25
PCIe_2_C_RXN	PCIe RX- of Lane C Link 2	Host connect	27
PCIe_2_D_TXP	PCIe TX+ of Lane D Link 2	Host connect	30
PCIe_2_D_TXN	PCIe TX- of Lane D Link 2	Host connect	32
PCIe_2_D_RXP	PCIe RX+ of Lane D Link 2	Host connect	33
PCIe_2_D_RXN	PCIe RX- of Lane D Link 2	Host connect	35
PCIe_3_A_TXP	PCIe TX+ of Lane A Link 3	Host connect	38
PCIe_3_A_TXN	PCIe TX- of Lane A Link 3	Host connect	40
PCIe_3_A_RXP	PCIe RX+ of Lane A Link 3	Host connect	41
PCIe_3_A_RXN	PCIe RX- of Lane A Link 3	Host connect	43
PCIe_3_B_TXP	PCIe TX+ of Lane B Link 3	Host connect	46
PCIe_3_B_TXN	PCIe TX- of Lane B Link 3	Host connect	48
PCIe_3_B_RXP	PCIe RX+ of Lane B Link 3	Host connect	49
PCIe_3_B_RXN	PCIe RX- of Lane B Link 3	Host connect	51
PCIe_3_C_TXP	PCIe TX+ of Lane C Link 3	Host connect	54
PCIe_3_C_TXN	PCIe TX- of Lane C Link 3	Host connect	56
PCIe_3_C_RXP	PCIe RX+ of Lane C Link 3	Host connect	57
PCIe_3_C_RXN	PCIe RX- of Lane C Link 3	Host connect	59
PCIe_3_D_TXP	PCIe TX+ of Lane D Link 3	Host connect	62
PCIe_3_D_TXN	PCIe TX- of Lane D Link 3	Host connect	64
PCIe_3_D_RXP	PCIe RX+ of Lane D Link 3	Host connect	65
PCIe_3_D_RXN	PCIe RX- of Lane D Link 3	Host connect	67
Clk1_100M_PCIe_P	100MHz clk +	PCIe Clk	70
Clk1_100M_PCIe_N	100MHz clk -	PCIe Clk	72
RSVD	Reserved pins	Future use	1, 3, 73, 75, 77, 78, 79,

Signal Name	Signal Description	Purpose	Connector Location
			80
Ground	Ground		2, 4, 5, 7, 10, 12, 13, 15, 18, 20, 21, 23, 26, 28, 29, 31, 34, 36, 37, 39, 42, 44, 45, 47, 50, 52, 53, 55, 58, 60, 61, 63, 66, 68, 69, 71, 74, 76

Together, the two blade/mezzanine connectors support four x4 PCIe links to the Compute Blade that can be configured as shown Table 1-7.

Table 1-7: PCI Express Configurations

PCIe_0	PCIe_1	PCIe_2	PCIe_3
X4	X4	X4	X4
X8		X4	X4
X8		X8	
X16			

In addition, each x4 link could be configured into a x1 link by leaving the high-order lanes of that link unconnected. If not all links are utilized, then low order links **shall** be used first. For example, if only one x4 is utilized, then PCIe_0 **shall** be used.

For a pass-through PCIe implementation, no more that two x4 links may pass-through to the mid-plane Airmax connector.

Note that if only the primary blade/mezzanine connector is implemented, then the x16 configuration is not possible.

1.2.4 Power Pins

Power delivery to the mezzanine PBA is via 12V, 5V, and 3.3V signals from the Compute Blade. Table 1-8 lists the maximum current capability of each rail.

Table 1-8: Power Rails Maximum Current Draw

	12V	5V	3.3V	3.3V AUX
Current	2A	0.5A	2A	1.5A

The mezzanine PBA cannot exceed the maximum current on any of the rails, and it is limited to a maximum of 25W total. The Compute Blade **shall** provide the maximum current for all rails.

Power is delivered to the mezzanine PBA from the Compute Blade. Mezz_PRES_N (pin 98 on the primary connector) indicates to the Compute Blade that a mezzanine card is present.

The Compute Blade will power up the 3.3V AUX power (P3V3AUX, pins 13-16 on the primary connector) first, so the mezzanine PBA can communicate capabilities and management information prior to power up. The main rails **shall** be controlled on the Compute Blade such that they are not powered until after 3.3V AUX is stable. Typically, this is after the Compute Blade has received permission to power up from Chassis Management Module. Other than being powered up after 3.3V AUX, there are no sequence requirements for the main rails. Note: the channels to the mid-plane **shall not** be powered with AUX power.

In addition to the power delivery pins, there are SMB signals for mezzanine PBA management from the Compute Blade, and signals to drive LEDs, if supported, by a Compute Blade.

1.2.5 LED Control

A mezzanine PBA **shall** provide support for a pair of active low LVTTL LED control signals for each of the four Flexi-Interconnect links the mezzanine supports. These signals may optionally be used by a Compute Blade to drive LEDs.

An example of an implementation for an Ethernet link is as follows. Other implementations are possible.

Each of the active low LNK_LED_N signals (pins 19, 21, 23, and 25 on the primary connector) is asserted to indicate its corresponding link is established.

Each of the active low ACT_LED_N signals (pins 20, 22, 24, and 26 on the primary connector) is pulsed low in a manner such that an LED can visually blink during activity on its corresponding link.

To define the LED priority for single LED connections (per port) between a Mezzanine and a Compute Blade, the pin labeled ACT_LED_N **should** be connected first on both modules.

1.2.6 Mezzanine PBA to Mid-plane PBA Connector Pinout

As in the Compute Blade specification, a pair of differential signals (4 copper traces) is considered a channel. The mezzanine PBA to Flexi mid-plane interface supports two 4-channel links and two 1-channel links. Links are labeled numerically (0 through 3), while channels are labeled alphabetically (A through D). Additional pins are reserved for future use.

If connection redundancy is required, both interfaces in a pair **should** be used, and both interfaces must use the same technology. For example, for redundancy in the 4-channel link, both links **should** be used, and they must be the same technology. Identical technologies are only required in a pair of links, not across all four links.

An Airmax VS 96-pin connector **shall** be used with the pinout shown in Table 1-9 and Table 1-10.

Note: The mezzanine PBA to Flexi mid-plane signal directions are with respect to the mezzanine. For example, RX goes to the receiver on the mezzanine.

Table 1-9: Mezzanine PBA to Flexi Mid-plane Connector

	A	B	C	D	E	F	G	H	I	J	K	L	
8	GND	HSC_0_A_TXP	HSC_0_A_TXN	GND	HSC_1_A_TXP	HSC_1_A_TXN	GND	HSC_2_A_TXP	HSC_2_A_TXN	GND	HSC_3_A_TXP	HSC_3_A_TXN	8
7	HSC_0_A_RXP	HSC_0_A_RXN	GND	HSC_1_A_RXP	HSC_1_A_RXN	GND	HSC_2_A_RXP	HSC_2_A_RXN	GND	HSC_3_A_RXP	HSC_3_A_RXN	GND	7
6	GND	HSC_0_B_TXP	HSC_0_B_TXN	GND	HSC_1_B_TXP	HSC_1_B_TXN	GND	RSVD	RSVD	GND	RSVD	RSVD	6
5	HSC_0_B_RXP	HSC_0_B_RXN	GND	HSC_1_B_RXP	HSC_1_B_RXN	GND	RSVD	RSVD	GND	RSVD	RSVD	GND	5
4	GND	HSC_0_C_TXP	HSC_0_C_TXN	GND	HSC_1_C_TXP	HSC_1_C_TXN	GND	RSVD	RSVD	GND	RSVD	RSVD	4
3	HSC_0_C_RXP	HSC_0_C_RXN	GND	HSC_1_C_RXP	HSC_1_C_RXN	GND	RSVD	RSVD	GND	RSVD	RSVD	GND	3
2	GND	HSC_0_D_TXP	HSC_0_D_TXN	GND	HSC_1_D_TXP	HSC_1_D_TXN	GND	RSVD	RSVD	GND	RSVD	RSVD	2
1	HSC_0_D_RXP	HSC_0_D_RXN	GND	HSC_1_D_RXP	HSC_1_D_RXN	GND	RSVD	RSVD	GND	RSVD	RSVD	GND	1
	A	B	C	D	E	F	G	H	I	J	K	L	

Table 1-10: Mezzanine PBA to Flexi Mid-plane Signal Description

Signal Name	Signal Description	Purpose	Connector Location
HSC_0_A_TXP	HSC TX+ of Channel A Link 0	High Speed Midplane Connect	B8
HSC_0_A_TXN	HSC TX- of Channel A Link 0	High Speed Midplane Connect	C8
HSC_0_A_RXP	HSC RX+ of Channel A Link 0	High Speed Midplane Connect	A7
HSC_0_A_RXN	HSC RX- of Channel A Link 0	High Speed Midplane Connect	B7
HSC_0_B_TXP	HSC TX+ of Channel B Link 0	High Speed Midplane Connect	B6
HSC_0_B_TXN	HSC TX- of Channel B Link 0	High Speed Midplane Connect	C6
HSC_0_B_RXP	HSC RX+ of Channel B Link 0	High Speed Midplane Connect	A5
HSC_0_B_RXN	HSC RX- of Channel B Link 0	High Speed Midplane Connect	B5
HSC_0_C_TXP	HSC TX+ of Channel C Link 0	High Speed Midplane Connect	B4
HSC_0_C_TXN	HSC TX- of Channel C Link 0	High Speed Midplane Connect	C4
HSC_0_C_RXP	HSC RX+ of Channel C Link 0	High Speed Midplane Connect	A3
HSC_0_C_RXN	HSC RX- of Channel C Link 0	High Speed Midplane Connect	B3
HSC_0_D_TXP	HSC TX+ of Channel D Link 0	High Speed Midplane Connect	B2
HSC_0_D_TXN	HSC TX- of Channel D Link 0	High Speed Midplane Connect	C2
HSC_0_D_RXP	HSC RX+ of Channel D Link 0	High Speed Midplane Connect	A1
HSC_0_D_RXN	HSC RX- of Channel D Link 0	High Speed Midplane Connect	B1
HSC_1_A_TXP	HSC TX+ of Channel A Link 1	High Speed Midplane Connect	E8
HSC_1_A_TXN	HSC TX- of Channel A Link 1	High Speed Midplane Connect	F8
HSC_1_A_RXP	HSC RX+ of Channel A Link 1	High Speed Midplane Connect	D7
HSC_1_A_RXN	HSC RX- of Channel A Link 1	High Speed Midplane Connect	E7
HSC_1_B_TXP	HSC TX+ of Channel B Link 1	High Speed Midplane Connect	E6

Signal Name	Signal Description	Purpose	Connector Location
HSC_1_B_TXN	HSC TX- of Channel B Link 1	High Speed Midplane Connect	F6
HSC_1_B_RXP	HSC RX+ of Channel B Link 1	High Speed Midplane Connect	D5
HSC_1_B_RXN	HSC RX- of Channel B Link 1	High Speed Midplane Connect	E5
HSC_1_C_TXP	HSC TX+ of Channel C Link 1	High Speed Midplane Connect	E4
HSC_1_C_TXN	HSC TX- of Channel C Link 1	High Speed Midplane Connect	F4
HSC_1_C_RXP	HSC RX+ of Channel C Link 1	High Speed Midplane Connect	D3
HSC_1_C_RXN	HSC RX- of Channel C Link 1	High Speed Midplane Connect	E3
HSC_1_D_TXP	HSC TX+ of Channel D Link 1	High Speed Midplane Connect	E2
HSC_1_D_TXN	HSC TX- of Channel D Link 1	High Speed Midplane Connect	F2
HSC_1_D_RXP	HSC RX+ of Channel D Link 1	High Speed Midplane Connect	D1
HSC_1_D_RXN	HSC RX- of Channel D Link 1	High Speed Midplane Connect	E1
HSC_2_A_TXP	HSC TX+ of Channel A Link 2	High Speed Midplane Connect	H8
HSC_2_A_TXN	HSC TX- of Channel A Link 2	High Speed Midplane Connect	I8
HSC_2_A_RXP	HSC RX+ of Channel A Link 2	High Speed Midplane Connect	G7
HSC_2_A_RXN	HSC RX- of Channel A Link 2	High Speed Midplane Connect	H7
HSC_3_A_TXP	HSC TX+ of Channel A Link 3	High Speed Midplane Connect	K8
HSC_3_A_TXN	HSC TX- of Channel A Link 3	High Speed Midplane Connect	L8
HSC_3_A_RXP	HSC RX+ of Channel A Link 3	High Speed Midplane Connect	J7
HSC_3_A_RXN	HSC RX- of Channel A Link 3	High Speed Midplane Connect	K7

1.2.7 Mezzanine PBA Management Features

The mezzanine PBA **shall** support intelligent management through the SMB bus. Upon blade power up and before the I/O components are activated, the Chassis Management Module (via the blade management controller) reads the mezzanine PBA configuration, status, and FRU information. This information **shall** be stored on the mezzanine PBA in non-volatile storage at a specific offset. See the Compute Blade Mezzanine Management section for details.

The 3.3V AUX power **shall** always be present. It provides a means for the system management software to communicate with the mezzanine PBA. The mezzanine PBA **shall** provide a means for supporting AUX power so this method of power up can be supported.

1.2.8 Mezzanine PBA Interfaces to the Flexi Mid-plane Interconnect

A mezzanine PBA must implement connectivity that is compatible with both redundant and non-redundant chassis Flexi-Channel implementations. The connection to the switch in a chassis, however, is implementation specific. For example, either redundant (connected to separate switches) or non–redundant, aggregated (connected to two ports of one switch), determined by the OEMs.

If a mezzanine PBA supports single-channel technology that is supported by the SSI 4x High Speed Switch Modules (e.g., 10GBASE-KR), it **shall** implement support for Link 0 Channel A and Link 1 Channel A.

If the mezzanine PBA supports single-channel technology that is supported by the SSI 1x Low Speed Switch Module (e.g. 1000BASE -KX or Fibre Channel), then it **shall** implement support for Link 2 Channel A and Link 3 Channel A.

If a mezzanine PBA supports quad-channel technology (e.g., 10GBASE-KX4 or 8Gb PCIe Gen1), then both Link 0 and Link 1 **shall** be supported.

Note that Links 2 and 3 only support single-channel technology.

Combining x4 links to create a x8 link configuration is not supported.

1.2.8.1 Ethernet Configuration

If the mezzanine PBA connects to the chassis Ethernet fabric via full-link, 4-channel 10GBASE-KX4, or single-channel 10GBASE-KR connections, the connection **shall** support both Link 0 and Link 1. These implement a redundant pair of interconnections and are connected to separate switches.

If the mezzanine PBA connects to the Ethernet fabric using single-channel 1000BASE-KX configurations, it **shall** support both Link 2 and Link 3. These implement a redundant pair of interconnections.

The mezzanine configuration **shall** support the IEEE 802.3ap standard and **shall** support the auto-negotiation and detection methods specified in the standard.

1.2.8.2 Supported Channel Configuration

Table 1-11 shows the grouping of the serial differential pairs for single-channel and multi-channel link technologies, such as Fibre Channel and Infiniband. Note: Links 0 and 1 are intended to be routed through the midplane to SSI 4x High Speed Switch Modules. Links 2 and 3 are intended to be routed through the mid-plane to SSI 1x Low Speed Switch Modules.

Table 1-11: Multi-lane Serial Link Configuration

<table>
<tr><th colspan="4">Mezzanine to Fabric Interconnect Link 0
(Channels A-D)</th><th colspan="4">Mezzanine to Fabric Interconnect Link 1
(Channels A-D)</th></tr>
<tr><td>ChA</td><td>ChB</td><td>ChC</td><td>ChD</td><td>CHA</td><td>ChB</td><td>ChC</td><td>CHD</td></tr>
<tr><td>1X Serial Diff Pair</td><td colspan="3">Not Connected</td><td>1x Serial Diff Pair</td><td colspan="3">Not Connected</td></tr>
<tr><td colspan="2">X2 Serial Diff Pairs</td><td colspan="2">Not Connected</td><td colspan="2">X2 Serial Diff Pairs</td><td colspan="2">Not Connected</td></tr>
<tr><td colspan="4">X4 Serial Diff Pairs</td><td colspan="4">X4 Serial Diff Pairs</td></tr>
<tr><td colspan="8">X8 configuration is not supported.</td></tr>
</table>

2 Compute Blade Mezzanine Management

This specification defines two mezzanine management models – unintelligent and intelligent. These models refer to the type of management support on the mezzanine PBA itself. A mezzanine PBA **shall** support one of these management models. Unmanaged mezzanine PBAs are not allowed.

The mezzanine PBA is intended to supply additional network connectivity to the blade. However, there is no support for a side-band management connection, such as TCO or RMII, to the Compute Blade management subsystem. Therefore, a mezzanine cannot be used to support a redundant Compute Blade management Ethernet link.

2.1 IPMI FRU Inventory Support

The *IPMI v2.0* specification defines Field Replaceable Unit (FRU) Information commands and formats for querying and retrieving Inventory information from an IPMI capable entity. The *Intelligent Platform Management FRU Information Storage Definition v1.0 r1.1* specification defines the common format and use of the FRU Information using IPMI.

Independent of the mezzanine management model, the mezzanine PBA **shall** support IPMI FRU Inventory data which **shall**, minimally, contain the following areas:

- Common Header Area: This area contains pointers into the other IPMI FRU data areas.
- Board Info Area: The *Product Name*, *Manufacturer*, and *Serial Number* fields **shall** be supported.
- MultiRecord Area: This area **shall** contain a Mezzanine Specification Compliance record, a Module Current Draw record (defined below) and Sensor Definition and Signal Interconnect records, as appropriate.

All area and record checksums **shall** be correctly maintained.

Other fields and records may be supported at the manufacturer's discretion.

2.2 Specification Compliance

In order for the Compute Blade management controller to determine the capabilities of the mezzanine and overall Open Blade Mezzanine specification compliance, the mezzanine's IPMI FRU Inventory MultiRecord area **shall** contain a *Mezzanine Specification Compliance* record. It indicates the specific mezzanine specification the mezzanine PBA **shall** adhere to as well as provides a set of option indications (currently unused) for mezzanine implementation options. Table 2-1 defines the Mezzanine Specification Compliance record.

Table 2-1: Mezzanine PBA Specification Compliance Record

Offset	Length	Definition
0	1	*Record Type ID* – C0h (OEM) **shall** be used
1	1	End of List/Version [7] - End of list. Set to 1 for the last record [6:4] - Reserved, write as 0h [3:0] - Record format version = 2h for this definition
2	1	*Record Length* = 0Eh
3	1	*Record Checksum* – holds the zero checksum for the record
4	1	*Header Checksum* – holds the zero checksum for the header
5-7	3	*Manufacturer ID* **– LS byte first.** Write as the three byte ID assigned to SSI – 7244h for this specification (44,72,00h in the record in offset 5,6,7)
8	1	*Open Blade Record ID – Mezzanine Specification Compliance Record* = 23h
9	1	*Record Format Version* – version for this record type = 00h
10	1	*Mezzanine Specification Major Revision Number* 1 for this specification
11	1	*Mezzanine Specification Minor Revision Number* 0 for this specification
12-13	2	*Mezzanine Options* This is a 2 byte bitmask indicating the specification options implemented by the mezzanine. Currently there are no options defined. [15:0] - Reserved, must be 0

2.3 Power Budget Support

In order to allow the associated Compute Blade management subsystem to correctly estimate the blade's power budget, the mezzanine's IPMI FRU Inventory MultiRecord area **shall** contain a *Module Current Draw* record which defines the maximum current draw of the mezzanine on the three (12v, 5v, and 3.3v) power rails. Table 2-2 defines the Module Current Draw record.

Table 2-2: Module Current Draw Record

Offset	Length	Definition
0	1	*Record Type ID* – C0h (OEM) **shall** be used
1	1	*End of List/Version* [7] - End of list. Set to 1 for the last record [6:4] - Reserved, write as 0h [3:0] - Record format version = 2h for this definition
2	1	*Record Length* = 0Dh
3	1	*Record Checksum* – holds the zero checksum for the record
4	1	*Header Checksum* – holds the zero checksum for the header
5-7	3	*Manufacturer ID* **– LS byte first.** Write as the three byte ID assigned to SSI – 7244h for this specification (44,72,00h in the record in offset 5,6,7)
8	1	*Open Blade Record ID – Module Current Draw Record* = 20h
9	1	*Record Format Version* – version for this record type = 00h
10	1	*12v current draw* – in tenths of an Ampere.
11	1	*5v current draw* – in tenths of an Ampere.
12	1	*3.3v current draw* – in tenths of an Ampere.

2.4 Signal Interconnect Support

A mezzanine PBA provides connectivity between the Compute Blade and the Flexi mid-plane connectors. It is a management task to verify that the Compute Blade connections are compatible with the devices connected to them. To that end, a mezzanine IPMI FRU Info area contains a MultiRecord area *Signal Interconnect* record that associates interconnect technology with the Compute Blade interconnect channels driven by the mezzanine circuitry.

Link 0 Channels A-D, and Link 1 Channels A-D correspond to Channels 11-14 and 15-18 respectively in the Compute Blade *Get Channel State* and *Set Channel State* commands. Table 2-3 defines the Signal Interconnect record. defines the Technology–specific Support data.

Table 2-3: Module Signal Interconnect Record

Offset	Length	Definition
0	1	*Record Type ID* – C0h (OEM) **shall** be used
1	1	*End of List/Version* [7] - End of list. Set to 1 for the last record [6:4] - Reserved, write as 0h [3:0] - Record format version = 2h for this definition
2	1	*Record Length* = 1Bh
3	1	*Record Checksum* – holds the zero checksum for the record.
4	1	*Header Checksum* – holds the zero checksum for the header
5-7	3	*Manufacturer ID* **– LS byte first.** Write as the three byte ID assigned to SSI – 7244h for this specification (44,72,00 in the record in offset 6,7,8)
8	1	*Open Blade Record ID – Module Signal Interconnect Record* = 22h
9	1	*Record Format Version* – version for this record type = 00h
10	1	*Link 0 Technology* 00h = Not Implemented 01h = PCI Express 02h = Ethernet 03h = Fibre Channel 04h = Infiniband® 05h = SAS 06h = SATA 07h - FFh = Reserved
11	1	*Link 0 Technology Detail* 00h = Not Implemented (See Table 2-4 for specific values)
12	1	*Link 0 Width* Number of Lanes used for the link 00h = Not Implemented 01h = x1

Offset	Length	Definition
		02h = x2 03h = x3 04h = x4
13	1	*Link 0 Reserved*
14	1	*Link 1 Technology*
15	1	*Link 1 Technology Detail*
16	1	*Link 1 Width*
17	1	*Link 1 Reserved*
18	1	*Link 2 Technology*
19	1	*Link 2 Technology Detail*
20	1	*Link 2 Width*
21	1	*Link 2 Reserved*
22	1	*Link 3 Technology*
23	1	*Link 3 Technology Detail*
24	1	*Link 3 Width*
25	1	*Link 3 Reserved*
26	1	RSVD

Table 2-4: Technology-specific Support

Technology	Value	Type
PCI Express (01h)	01h	PCIe Gen1
	02h	PCIe Gen2
	03h-FFh	reserved for future use
Ethernet (02h)	01h	1000BASEKX
	02h	10GBASEKR
	03h	10GBASEKX4
	04h-FFh	reserved for future use
FibreChannel (03h)	01h	2GFC
	02h	4GFC
	03h	8GFC
	04h-FFh	reserved for future use
Infiniband (04h)	01h	Infiniband SDR
	02h	Infiniband DDR
	03h	Infiniband QDR
	04h-FFh	reserved for future use
SAS (05h)	01h	3Gb
	02h	6Gb
	03h-FFh	reserved for future use
SATA (06h)	01h	1.5Gb
	02h	3Gb
	03h	6Gb
	04h-FFh	reserved for future use

Table 2-5: Hardware Address Table

Offset	Length	Definition
0	1	*Record Type ID* – C0h (OEM) **shall** be used
1	1	*End of List/Version* [7] - End of list. Set to 1 for the last record [6:4] - Reserved, write as 0h [3:0] - Record format version = 2h for this definition
2	1	*Record Length* = 0Bh + N (length of address table)
3	1	*Record Checksum* – holds the zero checksum for the record.
4	1	*Header Checksum* – holds the zero checksum for the header

<table>
<tr><th>Offset</th><th>Length</th><th>Definition</th></tr>
<tr><td>5-7</td><td>3</td><td>Manufacturer ID – LS byte first. Write as the three byte ID assigned to SSI – 7244h for this specification (447200h in the record in offset 6,7,8)</td></tr>
<tr><td>8</td><td>1</td><td>Open Blade Record ID – Address Table = 24h</td></tr>
<tr><td>9</td><td>1</td><td>Record Format Version – version for this record type = 00h</td></tr>
<tr><td>10</td><td>1</td><td>Number of Addresses</td></tr>
<tr><td>11</td><td>N</td><td>Address Table Entries
The intent of this table is to store hardware addresses associated with the links described in the Module Signal Interconnect Record. As an example, a dual Ethernet NIC might have 2 MAC addresses defined here (one for Link0 and one for Link1). If the device supports more than one address per link, then this table can be “one-to-many” by having multiple addresses assigned to the same Link ID.

<table>
<tr><th>Offset</th><th>Length</th><th>Definition</th></tr>
<tr><td>0</td><td>1</td><td>Link ID (0 – 3) as enumerated in the Module Signal Interconnect Record</td></tr>
<tr><td>1</td><td>1</td><td>Address Length (L)</td></tr>
<tr><td>2</td><td>L</td><td>Hardware Address
Entries shall be stored as a series of individual bytes, first byte first.</td></tr>
</table>
…. Additional Table Entries starting at offset 2+L+1</td></tr>
</table>

2.5 Management Subsystem Power Draw

The mezzanine PBA management subsystem **shall** draw no more than 1.5A from the 3.3V AUX supply.

2.6 Unintelligent Mezzanine

An unintelligent mezzanine PBA has no management controller. The management support it does have is implemented by I2C devices of particular types from a set specified in this section.

2.6.1 IPMI FRU Inventory Support

An unintelligent mezzanine PBA **shall** provide an I2C Serial EEPROM of type 24C01, 24C02 (or compatible) that is capable of 100 kHz operation at a minimum, and **shall** be inter-operable with 400 KHz devices.

This device **shall** be at I2C address A0h. The device **shall** be powered via management power (3.3V AUX). The device may be hardware write protected, although that obviates use for customer asset information.

2.6.2 Sensor Support

An unintelligent mezzanine PBA **shall** provide one temperature sensor. This **shall** be implemented by an I2C temperature sensor device from the list of supported devices as defined in Table 2-6. Any device selected **shall**:

- Be capable of 100 kHz I2C operation at a minimum.
- Be inter-operable with 400KHz devices.
- Be configured to answer at one of the indicated I2C addresses.
- Be powered by 3.3V AUX.

In addition, the mezzanine PBA **shall** have *a Module Sensor Definition Record* in the MultiRecord area of the IPMI FRU Inventory information.

Table 2-6: Module Sensor Definition Record

Offset	Length	Definition
0	1	*Record Type ID* – C0h (OEM) **shall** be used
1	1	*End of List/Version* [7] - End of list. Set to 1 for the last record [6:4] - Reserved, write as 0h [3:0] - Record format version = 2h for this definition
2	1	*Record Length* = 0Dh
3	1	*Record Checksum* – holds the zero checksum for the record
4	1	*Header Checksum* – holds the zero checksum for the header
5-7	3	*Manufacturer ID* **– LS byte first.** Write as the three byte ID assigned to SSI – 7244h for this specification (44,72,00 in the record in offset 5,6,7)
8	1	*Open Blade Record ID – Module Sensor Definition Record* = 21h
9	1	*Record Format Version* – version for this record type = 00h
10	1	*Sensor Device Type* 0 = Reserved 1 = DS1631/DS1731 (or functionally equivalent) 2 = LM75 (or functionally equivalent) 3-255 = Reserved
11	1	*I2C Address* This is the address for the device. It **shall** be one of 98h, 9Ah, 9Ch, or 9Dh.
12	1	*Upper Non-Critical Threshold* This is the Upper Non-critical Threshold value to be used by the blade management controller when implementing the IPMI sensor associated with this sensor device. In degrees C.
13	1	*Upper Critical Threshold* This is the Upper Critical Threshold value to be used by the blade

Offset	Length	Definition
		management controller when implementing the IPMI sensor associated with this sensor device. In degrees C.
14	1	*Upper Non-Recoverable Threshold* This is the Upper Non-recoverable Threshold value to be used by the blade management controller when implementing the IPMI sensor associated with this sensor device. In degrees C.
15	1	*Hysteresis* This is the Hysteresis value to be used by the blade management controller when implementing the IPMI sensor associated with this sensor device. In degrees C.

2.7 Intelligent Mezzanine

An intelligent mezzanine PBA has a management controller that implements IPMI management over an IPMB interface – a satellite controller, as defined by the IPMI specification.

The IPMB address of the management controller **shall** be C0h.

This controller **shall** implement the following IPMI features.

2.7.1 IPMI FRU Inventory Area Support

This includes the *Read FRU Data* and *Get FRU Info* commands. The *Write FRU Data* command may be supported.

2.7.2 Sensor Device Support

This includes the Device SDR commands: *Get Device SDR Info*, *Get Device SDR*, and *Reserve Device SDR Repository* and the *Get Sensor Reading* command.

The mezzanine PBA management controller is an Event Generator, forwarding sensor events to the Compute Blade management controller, and so must support the Platform Event Message (to send events to the Compute Blade management controller), *Set Event Receiver*, and *Get Event Receiver* commands. The Event Receiver address **shall** default to 20h.

The mezzanine PBA management controller **shall** implement at least one temperature sensor.

The management controller's SDR Repository **shall** contain a *Management Controller Locator Record* and an SDR for the required temperature sensor.

2.7.3 IPMI Global Command Support

As required by the IPMI specification of management controllers, the mezzanine PBA management controller **shall** implement the *Get Device ID*, *Get Self Test Results*, and *Get Device GUID* commands. Table 2-7 lists the required IPMI commands.

Table 2-7: Mezzanine PBA Management Controller Required IPMI Commands

IPMI Command	NetFn	CMD Number
Get Device ID	App	01h
Get Self Test Results	App	04h
Get Device GUID	App	08h
Set Event Receiver	Sensor/Event	00h
Get Event Receiver	Sensor/Event	01h
Platform Event Msg	Sensor/Event	02h
Get Device SDR Info	Sensor/Event	20h
Get Device SDR	Sensor/Event	21h
Reserve Device SDR Repository	Sensor/Event	22h
Get Sensor Reading	Sensor/Event	2Dh
Get FRU Inventory Info	Storage	10h
Read FRU Data	Storage	11h

3 High Speed I/O Signal Specification

The mezzanine PBA to the Flexi-Interconnect signal specification is the same as the Compute Blade to the Primary Interconnect with the additional requirement that all Flexi-Interconnect signals to be between 2 and 5 inches on the mezzanine board.

See the Compute Blade Specification for details.

The PCIe signaling **shall** follow the PCIe Express Base Specification Revision 2.0.

Midplane connector pin skew length compensation **shall** be handled on the midplane. The mezzanine trace design **should not** compensate for the midplane connector signal pin skew length mismatch.

4 Compute Blade Mezzanine Thermal Management

4.1 Introduction

Airflow requirements and limitations for the mezzanine PBA are determined by several factors, including but not limited to the following: total heat dissipated, component selection, component density, component location, heat dissipation per device, acoustic targets, and practical limits of forced air cooling technology.

No general specification can dictate exactly how to adequately cool an unknown design, but general guidelines are presented herein. This specification is intended to enable air cooling of a wider variety of forward-looking mezzanine PBA designs.

4.2 Equipment Environment Specifications

See Table 2.1 on page 10 of Reference ASHRAE “Thermal Guidelines for Data Processing Environments” ISBN 1-931862-43-5, Class 1 for blade operating temperatures, humidity, and elevations. [302]

4.3 Mezzanine Airflow / Cooling Requirements

Refer to the Compute Blade’s thermal section for details.

5 Product Regulations Compliance

The Mezzanine product **shall** meet regulatory requirements as governed by specific country regulations.