OPEN

Compute Project

Inspur Mission Bay 3OU 8 GPU JBOG Syste Specification V0.1

Authors:
Inspur Mission Bay team

Table of Contents

1: License

OWF Contributor License Agreement 1.0 - Copyright and Patent

Open Web Foundation

Contributor License Agreement (CLA 1.0)

(Patent and Copyright Grants)

1. The Purpose of this Contributor License Agreement. This CLA sets forth the terms under which I will participate in and contribute to the development of the Specification. Capitalized terms are defined in the CLA's last section.

2. Copyrights.

2.1. Copyright Grant. I grant to you a perpetual (for the duration of the applicable copyright), worldwide, non-exclusive, no-charge, royalty-free, copyright license, without any obligation for accounting to me, to reproduce, prepare derivative works of, publicly display, publicly perform, sublicense, distribute, and implement any Contribution to the full extent of my copyright interest in the Contribution.

2.2. Attribution. As a condition of the copyright grant, you must include an attribution to the Specification in any derivative work you make based on the Specification. That attribution must include, at minimum, the Specification name and version number.

3. Patents.

3.1. Patent Non-Assert.

3.1.1. The Promise. I, on behalf of myself and my successors in interest and assigns, irrevocably promise not to assert my Granted Claims against you for your Permitted Uses, subject to the terms and conditions of Section 3.1. This is a personal promise directly from me to you, and you acknowledge as a condition of benefiting from it that no rights from me are received from suppliers, distributors. or otherwise in connection with this promise. This promise also applies to your Permitted Uses of any other specifications incorporating all required portions of the Specification.

3.1.2. Termination.

3.1.2.1. As a Result of Claims by You. All rights, grants, and promises made by me to you under this CLA are terminated if you file, maintain, or voluntarily participate in a lawsuit against me or any person or entity asserting that its Permitted Uses infringe any Granted Claims you would have had the right to enforce had you signed this CLA, unless that suit was in response to a corresponding suit first brought against you.

3.1.2.2. As a Result of Claims by a Related Entity of Mine. If a Related Entity of mine files, maintains, or voluntarily participates in a lawsuit asserting that a Permitted Use infringes any Granted Claims it would have had the right to enforce had it signed this CLA, then I relinquish any rights, grants, and promises I have received for the Specification from other signatories of this CLA, unless a) my promise to you was terminated pursuant to section 3.1.2.1, or b) that suit was in response to a corresponding suit first brought by you against the Related Entity.

3.1.3. Additional Conditions. This promise is not an assurance (i) that any of my copyrights or issued patent claims cover an implementation of the Specification or are enforceable or (ii) that an implementation of the Specification would not infringe intellectual property rights of any third party. Notwithstanding the personal nature of my promise, this promise is intended to be binding on any future owner, assignee or exclusive licensee who has been given the right to enforce any Granted Claims against third parties.

3.1.4. Bankruptcy. Solely for purposes of Section 365(n) of Title 11, United States Bankruptcy Code and any equivalent law in any foreign jurisdiction, this promise will be treated as if it were a license and you may elect to retain your rights under this promise if I (or any owner of any patents or patent applications referenced herein), as a debtor in possession, or a bankruptcy trustee, reject this non-assert.

3.2. Patent License Commitment. In addition to rights granted in 3.1, on behalf of me and my successors in interest and assigns, I agree to grant to you a no charge, royalty free license to my Granted Claims on reasonable and non-discriminatory terms, where such license applies only to those Granted Claims infringed by the implementation of my Contribution(s) alone or by combination of my Contribution(s) with the Specification, solely for your Permitted Uses.

4. No Other Rights. Except as specifically set forth in this CLA, no other express or implied patent, trademark, copyright, or other property rights are granted under this CLA, including by implication, waiver, or estoppel.

5. Limited Opt-Out. I may withdraw my Contribution by providing written notice of that withdrawal within 45 days of submitting that Contribution. Notice of a Contribution withdrawal must be made, at minimum, in writing using the same communication mechanisms that were used to submit the corresponding Contribution and must include the exact material being withdrawn. Upon providing such valid notice, any obligations I incurred under this CLA for that particular identified Contribution will be null and void.

6. Open Web Foundation Agreement ("OWFa") version 1.0 Execution. I acknowledge that the goal of this CLA is to develop a specification that will be subject to the OWFa version 1.0. While I have no legal obligation to execute the OWFa version 1.0 for any version of the specification being developed under this CLA, I agree that the selection and terms of the OWFa version 1.0 will not be subject to negotiation.

7. Antitrust Compliance. I acknowledge that I may compete with other participants, that I am under no obligation to implement the Specification, that each participant is free to develop competing technologies and standards, and that each party is free to license its patent rights to third parties, including for the purpose of enabling competing technologies and standards.

OPEN
Compute Project

8. Non-Circumvention. I agree that I will not intentionally take or willfully assist any third party to take any action for the purpose of circumventing my obligations under this CLA.

9. Representations, Warranties and Disclaimers. I represent and warrant that 1) I am legally entitled to grant the rights and promises set forth in this CLA and 2) I will not intentionally include any third party materials in any Contribution unless those materials are available under terms that do not conflict with this CLA. IN ALL OTHER RESPECTS MY CONTRIBUTIONS ARE PROVIDED "AS IS." The entire risk as to implementing or otherwise using the Contribution or the Specification is assumed by the implementer and user. Except as stated herein, I expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the Contribution or the Specification. 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 CLA, 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. All of my obligations under Section 3 regarding the transfer, successors in interest, or assignment of Granted Claims will be satisfied if I notify the transferee or assignee of any patent that I know contains Granted Claims of the obligations under Section 3. Nothing in this CLA requires me to undertake a patent search.

10. Definitions.

10.1. Bound Entities. "Bound Entities" means the entity listed below and any entities that the Bound Entity Controls.

10.2. CLA. "CLA" means this document, which sets forth the rights, grants, promises, limitations, conditions, obligations, and disclaimers made available for my Contributions to the particular Specification.

10.3. Contribution. "Contribution" means any original work of authorship, including any modifications or additions to an existing work, that I intentionally submit for inclusion in the Specification, which is included in the Specification. For the purposes of this definition, "submit" means any form of electronic, oral, or written communication for the purpose of discussing and improving the Specification, but excluding communication that I conspicuously designate in writing as not a contribution.

10.4. Control. "Control" means direct or indirect control of more than 50% of the voting power to elect directors of that corporation, or for any other entity, the power to direct management of such entity.

10.5. Granted Claims. "Granted Claims" are those patent claims that I own or control, including those patent claims I acquire or control after the Date below, that are infringed by Permitted Uses. Granted Claims include only those patent claims that are infringed by the implementation of any portions of the Specification where the Specification describes the functionality causing the infringement in detail and does not merely reference the functionality causing the infringement. Granted Claims under this CLA exclude those patent claims that would be infringed by an implementation of the Specification if my Contribution to that Specification were removed.

10.6. I, Me, or My. "I," "me," or "my" refers to the signatory below and its Bound Entities, if applicable.

10.7. Permitted Uses. "Permitted Uses" means making, using, selling, offering for sale, importing or distributing any implementation of the Specification 1) only to the

OPEN
Compute Project

extent it implements the Specification and 2) so long as all required portions of the Specification are implemented. Permitted Uses do not extend to any portion of an implementation that is not included in the Specification.

10.8. Related Entities. "Related Entities" means 1) any entity that Controls the Bound Entity ("Upstream Entity"), and 2) any other entity that is Controlled by an Upstream Entity that is not itself a Bound Entity.

10.9. Specification. "Specification" means the Specification identified below as of the date of my last Contribution.

10.10. You or Your. "You," "you," or "your" means any person or entity who exercises copyright or patent rights granted under this CLA, and any person or entity you Control.

Identify the Specification here:

Mission Bay 40U JBOG

If signing this CLA as an individual:

I acknowledge that, depending upon local law or contractual agreements, when I am employed by or acting on behalf of another entity, the promises I make relating to this Specification may actually be obligations of that other entity. In such a situation, I represent that I have been authorized by that entity to make these promises. I also understand that certain Specification projects may require additional private identifying information or certifications from me before they accept my Contributions.

Signed name

Print name

Email address

Date

If signing this CLA as a Bound Entity:

I certify that I am authorized to execute this agreement on behalf of the Bound Entity named below, and that all promises made herein relating to this Specification are commitments of the Bound Entity.

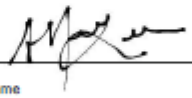

Signed name: Alfie Lew

Print name: ALFIE LEW

Email address: alfie@inspur.com

Date: 09/20/2018

Title: Sr. Solutions Architect

Bound Entity: Inspur Systems

Address

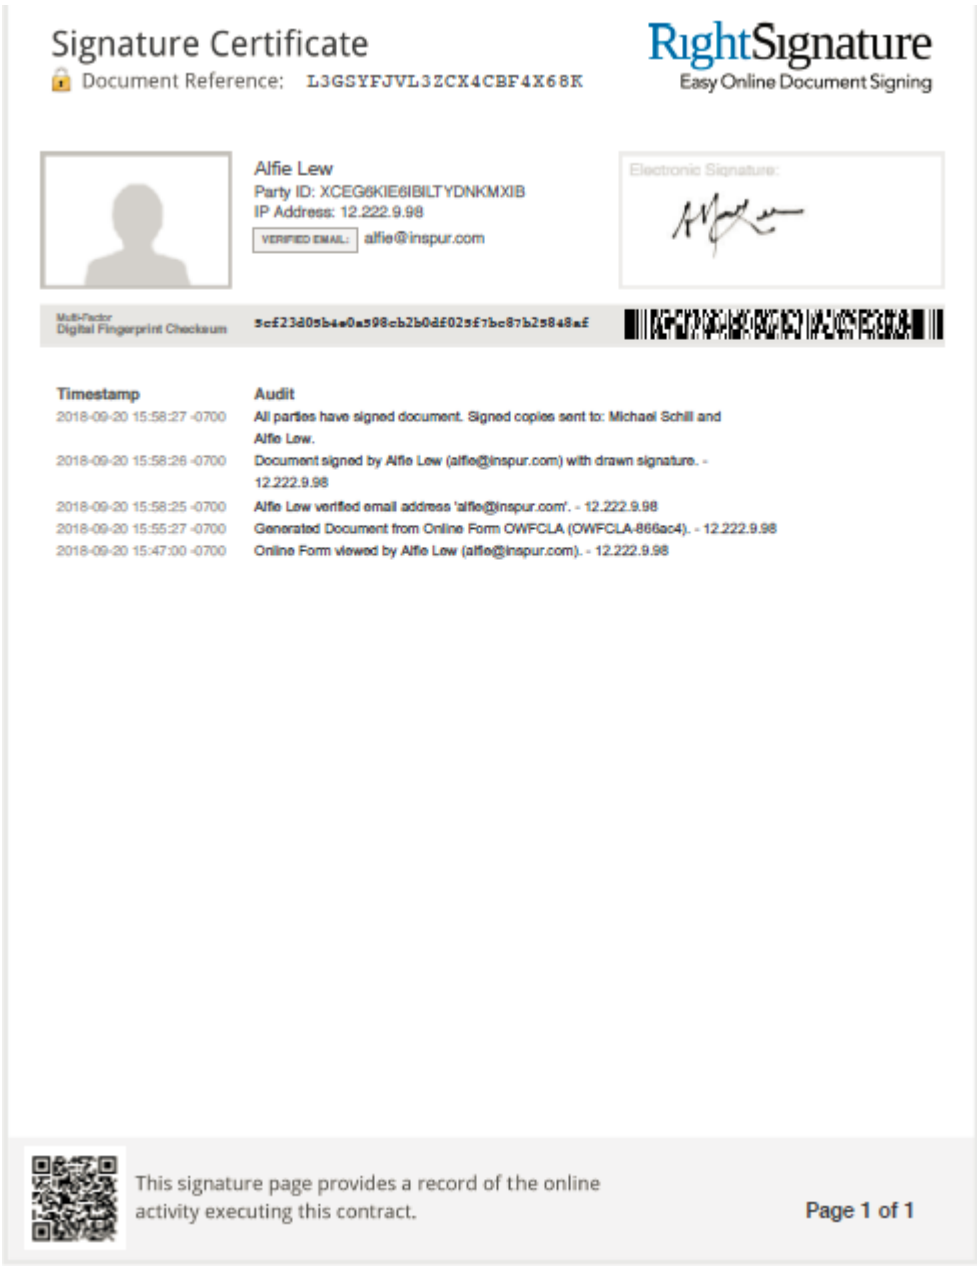

Signature Certificate

Document Reference: L3GSYFJVL3ZCX4CBF4X68K

RightSignature
Easy Online Document Signing

Alfie Lew
Party ID: XCEG6KIE6IBILTYDNKMXIB
IP Address: 12.222.9.98
VERIFIED EMAIL: alfie@inspur.com

Electronic Signature:

Multi-Factor Digital Fingerprint Checksum
5cf23d05b4e0a598cb2b0df025f7bc87b2584aef

Timestamp	Audit
2018-09-20 15:58:27 -0700	All parties have signed document. Signed copies sent to: Michael Schill and Alfie Lew.
2018-09-20 15:58:26 -0700	Document signed by Alfie Lew (alfie@inspur.com) with drawn signature. - 12.222.9.98
2018-09-20 15:58:25 -0700	Alfie Lew verified email address 'alfie@inspur.com'. - 12.222.9.98
2018-09-20 15:55:27 -0700	Generated Document from Online Form OWFCLA (OWFCLA-866ac4). - 12.222.9.98
2018-09-20 15:47:00 -0700	Online Form viewed by Alfie Lew (alfie@inspur.com). - 12.222.9.98

This signature page provides a record of the online activity executing this contract.

Page 1 of 1

2: Scope

This document defines the technical specifications for the Open Compute Project Mission Bay 8 GPU

3: Overview

The Mission Bay will be contents 8 GPU base board that supports NVIDIA SXM2 GPU to provide a flexible feature set between technology leadership and cost.

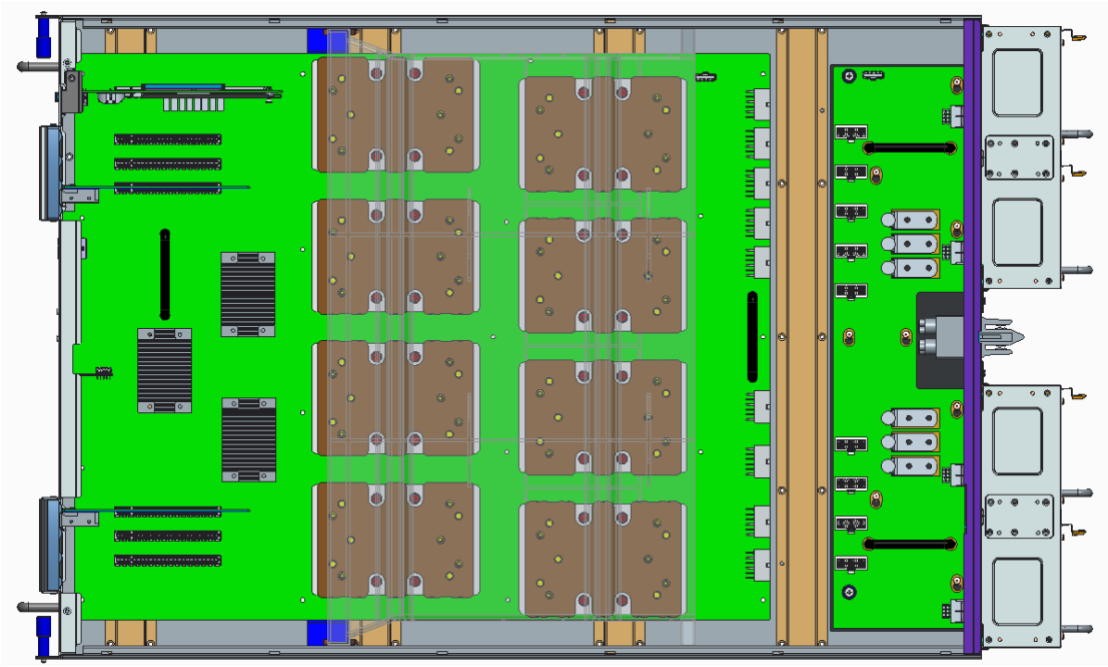

Figure 1: system top view

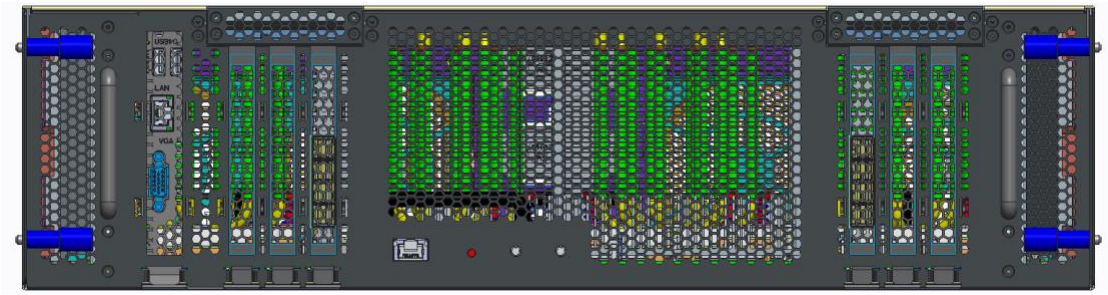

Figure 2: system front view

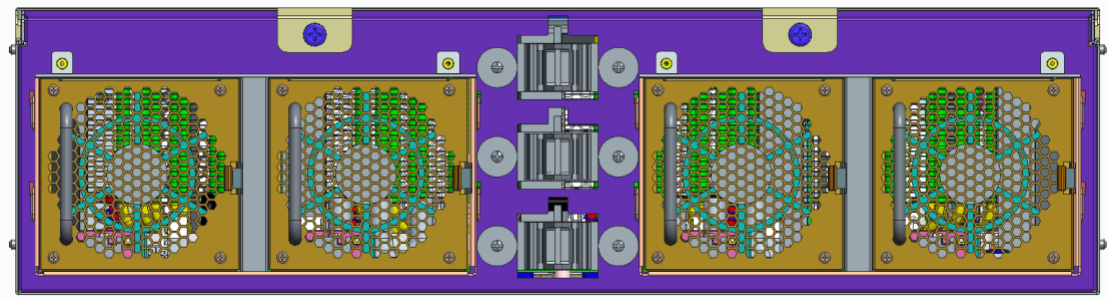

Figure 3: System rear view

4: Feature Requirements

4.1: PCB/Switch/GPU

Features		Description	Remark
PCB	**Form Factor**	No Standard	Inspur
	Layer	16	
	Color	Green	
	Board Size	450X590MM	
Switch	**Type**	PLX9797	
	Vendor	Broadcom	
	PCIE	PCIE 3.0	
	Lane	97 lanes	
	Thermal shutdown	Yes	
GPU	**Type**	SXM2 V100、P100	
	Vendor	INVIDIA	
	Port	2个PCIe3.0 x8 P100: 4个NVLink1.0 x8 V100: 6 个 NVLink2.0 x8	
	Thermal shutdown	Yes	

4.2 :Memory/Clock/Storage

Features	Description	Remark
Memory	No	
Generator	5P49V5901A750NLGI8	
Clock Buffer	9ZX21901BKLFT 9DB834AGILFT	

4.3: PCIE Slot

Features	Description	Remark
PCIEX16	PCIE A0，PEXA<16:31>，UP Stream	J40
PCIEX16	PCIE A1，PEXA<80:95>，Down Stream	J29
PCIEX16	PCIE A2，PEXC<80:95>，Down Stream	J30
PCIEX16	PCIE B0，PEXA<0:15>，UP Stream	J41
PCIEX16	PCIE B1，PEXA<64:79>，Down Stream	J28
PCIEX16	PCIE B2，PEXB<80:95>，Down Stream	J26
PCIEX24	CPU Board connector ,for fabric mode	J48

4.4: LAN

Features		Description	Remark
LAN1	**Vendor**	ASPEED	
	LAN Controller	AST2500	U134
	Transfers rate	10/100/1000Mb/s	

4.5: External I/O

Features		Description	Remark
External I/O	**1000M RJ45 with 1port**	X1	J56
	Power Button	X1	SW2
	UID Button	X1	SW1
	BMC Reset	X1	J123

4.6: Internal I/O

Features	Description (type, Q'ty)	Remark
PEX9797 EEPROM Socket	For Switch FW，x3	J44/J22/J21

PEX9797 I2C Connector	For swtich debug，x1	J33
PEX8608 I2C Connector	For 8608 debug，x1	J34
CPLD JTAG	For CPLD FW，x1	J63
GPU Connector	Two connector for one GPU，x16	J1/J2/,J3/J4,J5/J6, J7/J8,J9/J10,J11/J12/, J13/J14,J15/J16
GPU JTAG	X8	J17-J20/J23/J31/J25/J24
BMC FLASH Socket	For BMC FLASH，x1	J27
RTC Battery Socket	Battery，x1	J73
IPMB Connector	For IPMB，x1	J70
Power Connector	For Power，x9	J58-J62，J64-J67

4.7: Onboard header

Features		**Description(Q'ty)**	**Remark**
On-board header	**Fabric select**	X3	J37/J35/J32(x3)
	VS/Base Mode select	X1	J51(x3)
	MCU0/1 JTAG	Flash the MCU FW，x2	J50/J47(x5)
	MCU0/1 UART	For MCU debug，x2	J46/J45(x3)
	BMC Reload	X1	J98(x3)
	Phy disable	X1	J125(x3)
	Sharelink disable	X1	J126(x3)
	Video disable	X1	J91(x3)
	BMC Uart	X1	J104(x4)
	SYS Uart	X1	J53(x4)
	1278 I2C header	X1	J38(x3)
	9797 VR I2C header	X1	J36(x3)

5: Architecture Block Diagram

5.1: MB block diagram

The GPU board user 3 PEX9797s for PCIE fan out and PCIE management. It has 8 SXM2 GPUs mounted under the switch. It also use BMC for board management and PEX8608 for fabric usage. Its block diagram is as follows.

GPU Block Diagram

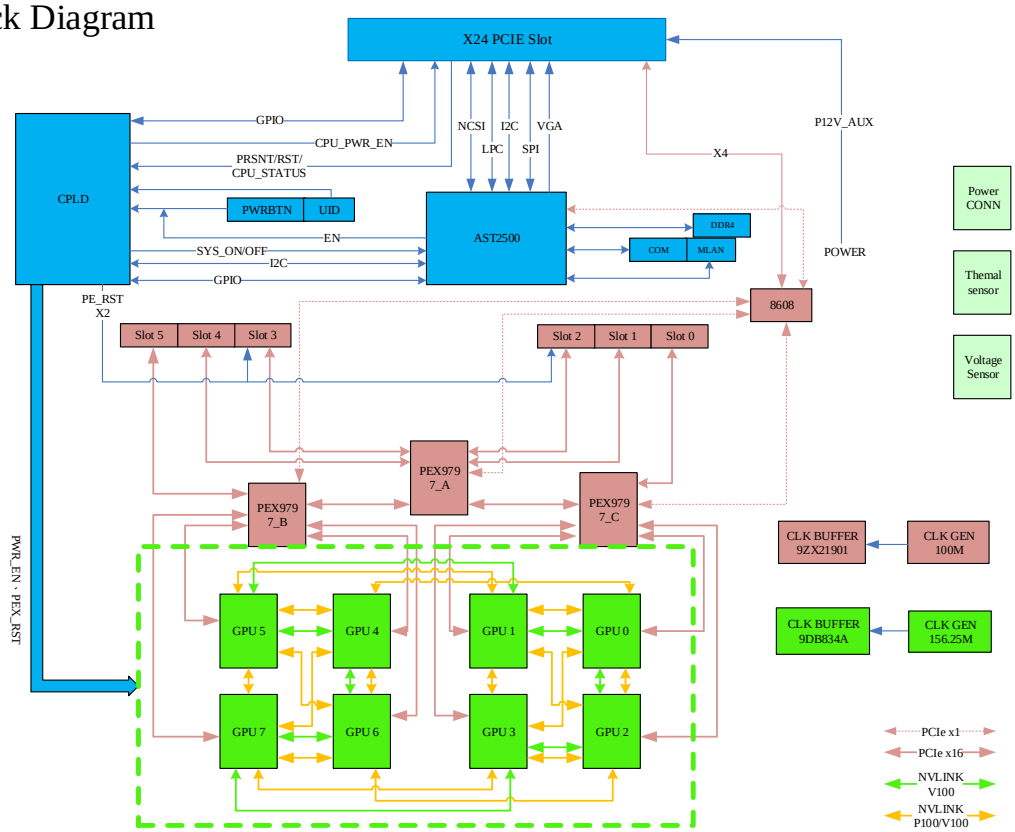

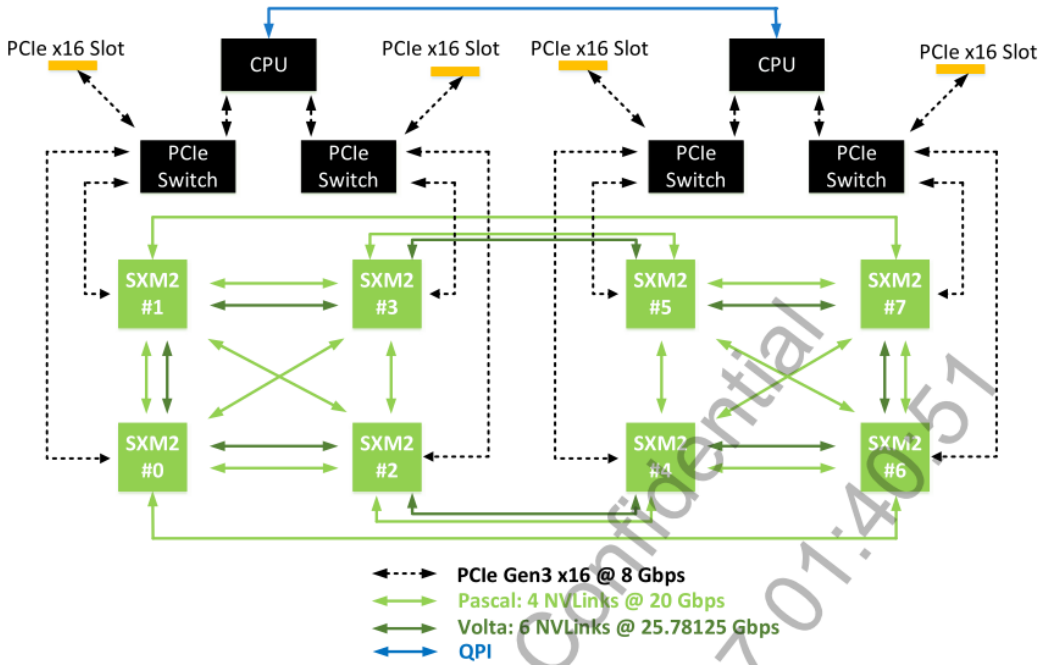

- PEX9797 x3, for PCIE fan out
- PEX8608 for PCIE management signal router
- SXM2 GPU is the PCIE device
- The MCU manages GPU power and sideband signals.
- The CPLD controls power on/off
- AST2500 for Box management
- 1 X24 GEN3 PCIE slot for management CPU board
- 2 x16 GEN3 PCIE slot for upstream to host
- 4 x16 GEN3 PCIE slot for downstream to endpoint
- 1 management LAN from BMC for box management
- 1 external uart for BMC debug
- 1 external uart for CPU debug

5.2: Chipset support

5.2.1 PEX9797 & GPU chipset:

PEX9797 is 96 lanes PCIE 3.0 PCIE switch. It has fanout mode and fabric mode.Fanout mode is the transparent bridge for host and endpoint.Fabric mode provide some advanced features such as PCIE management and Endpoint allocation etc,

In this GPU box，all of the three switches are configured as 6 x16 ports.The 9797_A port P0 and P4 are for upstream，others are for downstream.

The 8 GPUs are PCIe endpoints by PCIE x16 port. In the GPU box, the NVLink topology compatible with SXM2 V100 and P100.

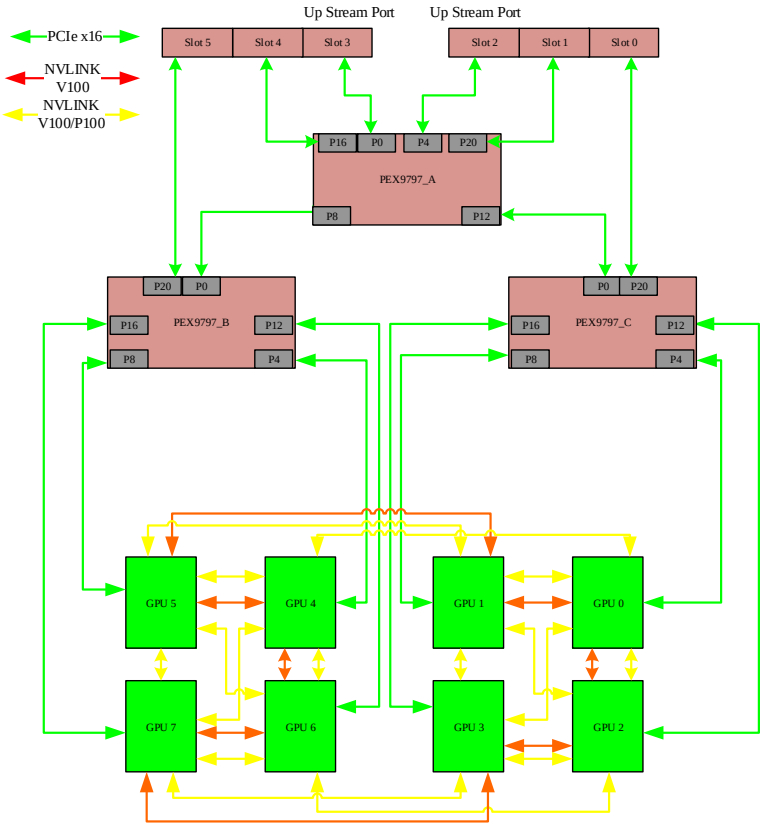

PCIE & NVLink topology

5.2.2 PEX9797 fabric mode

Fabric mode is for advanced research, the mode could support such as dual hosts, management PCIE endpoint and allocate endpoints from one to another. PEX9797 use the PCIE lane96 to connect management CPU(Avaton C2350, CPU Board) for the feature. For example, we can decide which gpu belongs to host0, and then change it to host1. The Avaton CPU is the management CPU for PEX9797, the management signal is route to PEX8608 first.

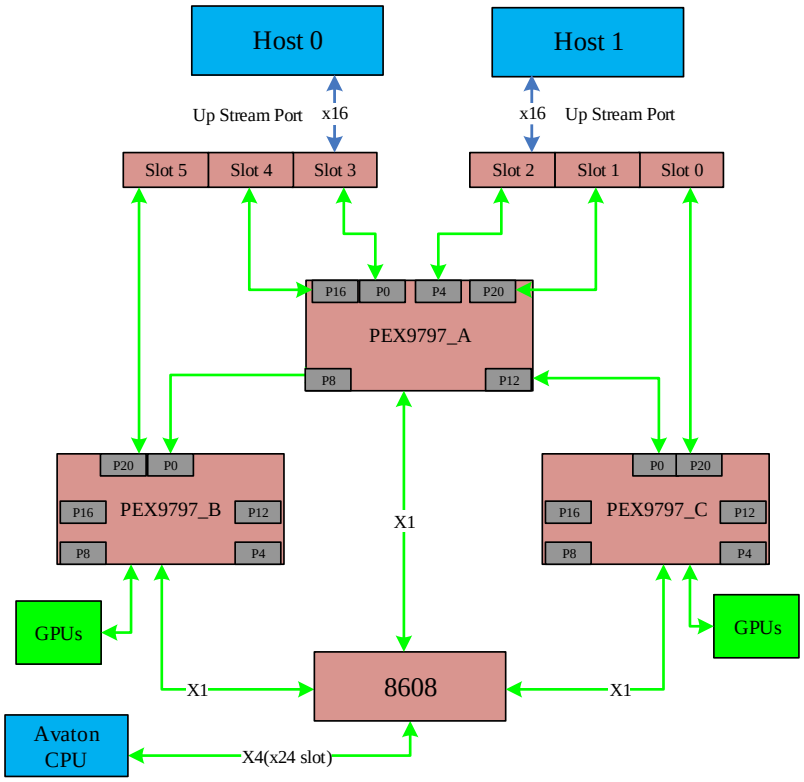

Fabric mode diagram

5.3 Clocking

The system use cfc clock generated local. All the PCIE switch and PCIE endpoint clocks come from the same clock buffer. The diagram as follows.

- 18X 100MHZ for PCIE
- 8X 156MHZ for GPU NVLink

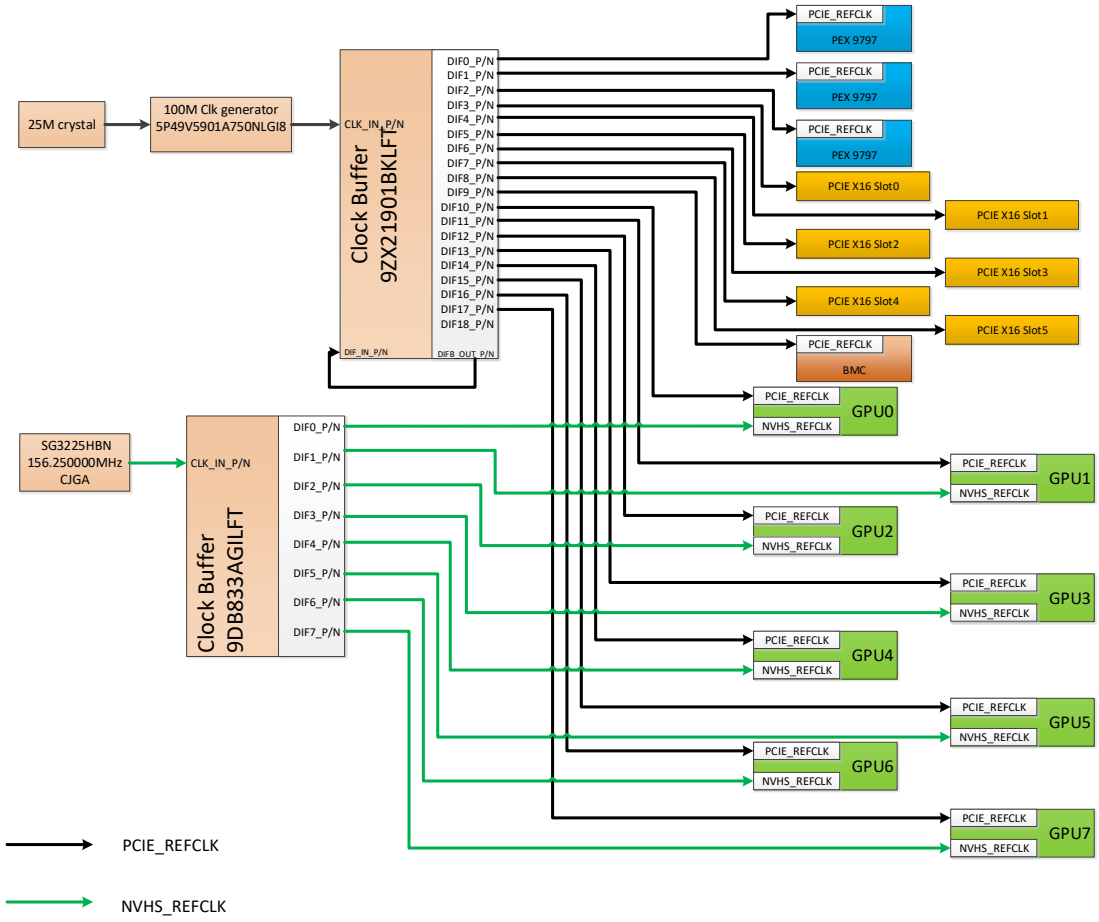

5.4: On Board Slot

Refer to 1.5 for details.

5.5: BMC

- ASPEED AST2500 chip with 4Gb DRAM
- Embedded NIC dedicated for management, 100/1000 Mbps port
- 1 COM ports support
- Embedded video (VGA) with 16MB video memory
- SMBUS 2.0 specification compliant
- Thermal sensor

5.6: I/O connector or onboard header

- 1 x1000M RJ45 connector for dedicated LAN

- 2 of 6 PCIe x16 slots for upstream, 4 ports for device.
- 1 x24 PICe 3.0 slot for management cpu.
- 1 x6pin I2C header for 9797 debug
- 1 x5pin JTAG header for MCU FW
- 1 x10pin JTAG header for CPLD FW
- 1 x30 pin IPMB connector for BMC communication, IPMB bus based on I2C interface

6: Environmental Requirements

This section provides regulatory and compliance information applicable to this system.

6.1 Product Regulatory Compliance Markings

This product shall be certified with the following Product Certification Markings for any deployments for their respective regions.

6.2 Operating environment

- Operating temperature: 5°C to 40°C (41°F to 104°F)
- Non-operating temperature: -40°C to 70°C (-40°F to 158°F)
- Operating relative humidity: 20% to 85%RH
- Non-operating relative humidity: 10% to 95%RH