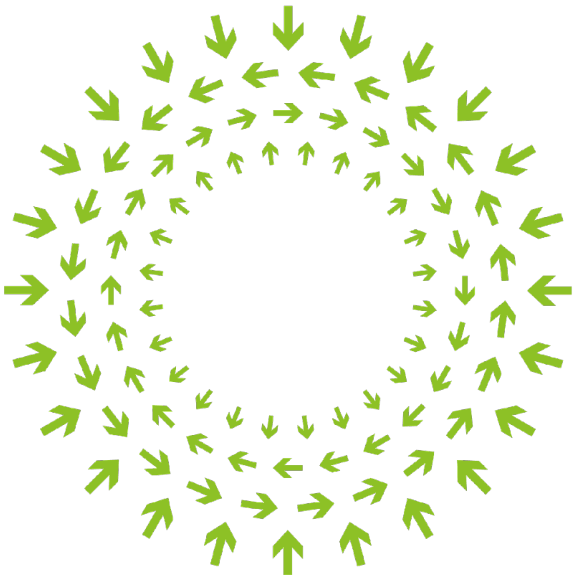

OPEN

Compute Project

Mt. Mitchell

Ampere@ CPU

Hardware Specification

V0.80

Author: **Leigh Chen**, Ampere Computing
Author: **Jayesh Shah**, Ampere Computing
Author: **Mike Yuan**, Inspur
Author: **Erin Duan**, Inspur

Table of Contents

1. License

PLEASE PICK EITHER THE OCP CLA OPTION OR THE OWF OPTION. ONLY ONE CAN BE USED. DELETE THE ONE NOT USED.

1.1. Open Web Foundation (OWF) CLA

Contributions to this Specification are made under the terms and conditions set forth in Open Web Foundation Modified Contributor License Agreement ("OWF CLA 1.0") ("Contribution License") by:

Inspur Corporation

Usage of this Specification is governed by the terms and conditions set forth in **Open Web Foundation Modified Final Specification Agreement ("OWFa 1.0") ("Specification License").**

You can review the applicable OWFa1.0 Specification License(s) referenced above by the contributors to this Specification on the OCP website at http://www.opencompute.org/participate/legal-documents/. For actual executed copies of either agreement, please contact OCP directly.

Notes:

1) The above license does not apply to the Appendix or Appendices. The information in the Appendix or Appendices is for reference only and non-normative in nature.

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL

OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING

NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

1.2. Acknowledgements

The Contributors of this Specification would like to acknowledge the following companies for their feedback:

Inspur Corporation

2. Compliance with OCP Tenets

Please describe how this Specification complies to the following OCP tenets. Compliance is required for at least three of the four tenets. The ideals behind open sourcing stipulate that everyone benefits when we share and work together. Any open source project is designed to promote sharing of design elements with peers and to help them understand and adopt those contributions. There is no purpose in sharing if all parties aren't aligned with that philosophy. The IC will look beyond the contribution for evidence that the contributor is aligned with this philosophy. The contributor actions, past and present, are evidence of alignment and conviction to all the tenets.

2.1. Openness

The Mt. Mitchell Design Specification exemplifies openness by providing a design for a server CPU which can be used to support numerous use cases ranging from general compute to storage to accelerators. This server design can be used as a standalone device in a standard 19” rack.

2.2. Efficiency

The form factor of the Mt. Mitchell server allows for better air flow and more efficient cooling for a higher powered CPU.

2.3. Impact

Mt. Mitchell is designed to dramatically reduce time to market for additional card and drive configurations. This server board can connect directly to custom risers and different NVMe drives on the backplane.

2.4. Scale

NA.

3. Revision Table

Date	Revision #	Author	Description
June 6, 2022	0.50	Leigh Chen	Initial Release
June 25, 2022	0.75	Leigh Chen	Updated with changes
September 9, 2022	0.80	Leigh Chen	Updated with changes

4. Scope

This specification describes the design of the Mt. Mitchell server board based on Ampere® processor.

5. Overview

All processor names mentioned in this specification refer to Ampere® processor. Details on Ampere® processors can be found at https://amperecomputing.com/.
This document describes a dual sockets or a single socket server design based on Ampere® processor, which is referred to hereinafter as Mt. Mitchell server motherboard.

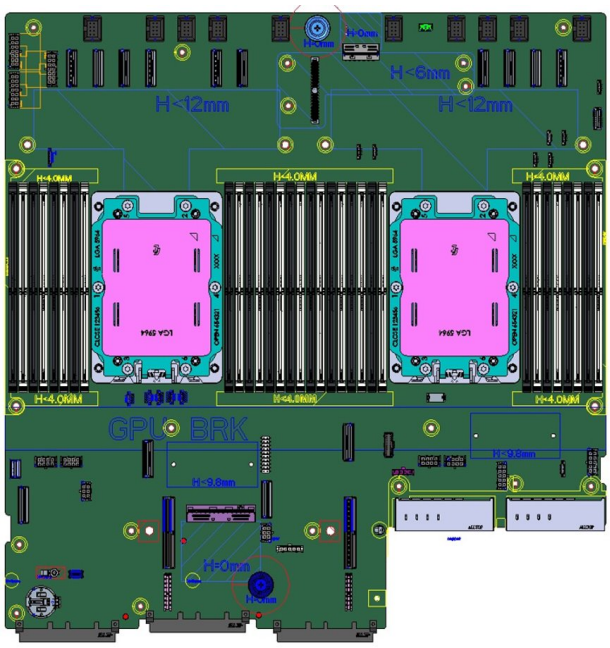

Figure 1: Mt. Mitchell 2 Sockets 32 DIMM Server Motherboard

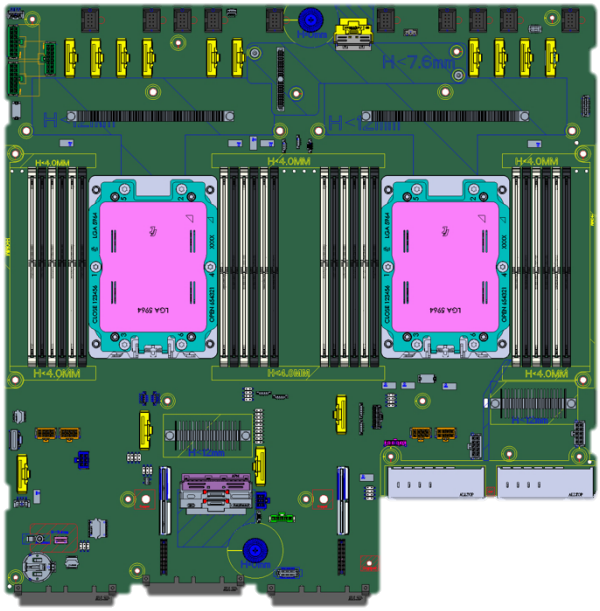

Figure 2: Mt. Mitchell 2 Sockets 24 DIMM Server Motherboard

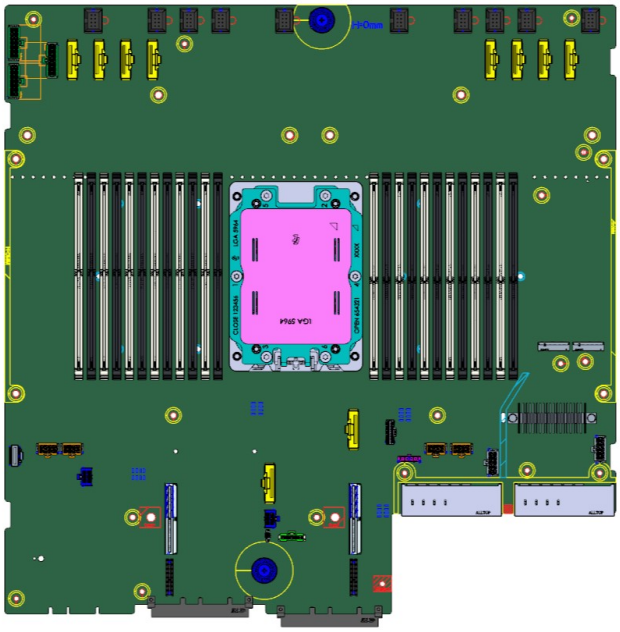

Figure 3: Mt. Mitchell 1 Socket 24 DIMM Server Motherboard

Mt. Mitchell server motherboard is designed with dual sockets, DDR5 DIMMs, 4C+ connector for x16 PCIE Gen5 OCPv3 NIC, riser connectors for PCIe riser cards, discrete ASPEED AST2600 BMC on DC-SCM 1.0, and two CRPS PSU connectors.

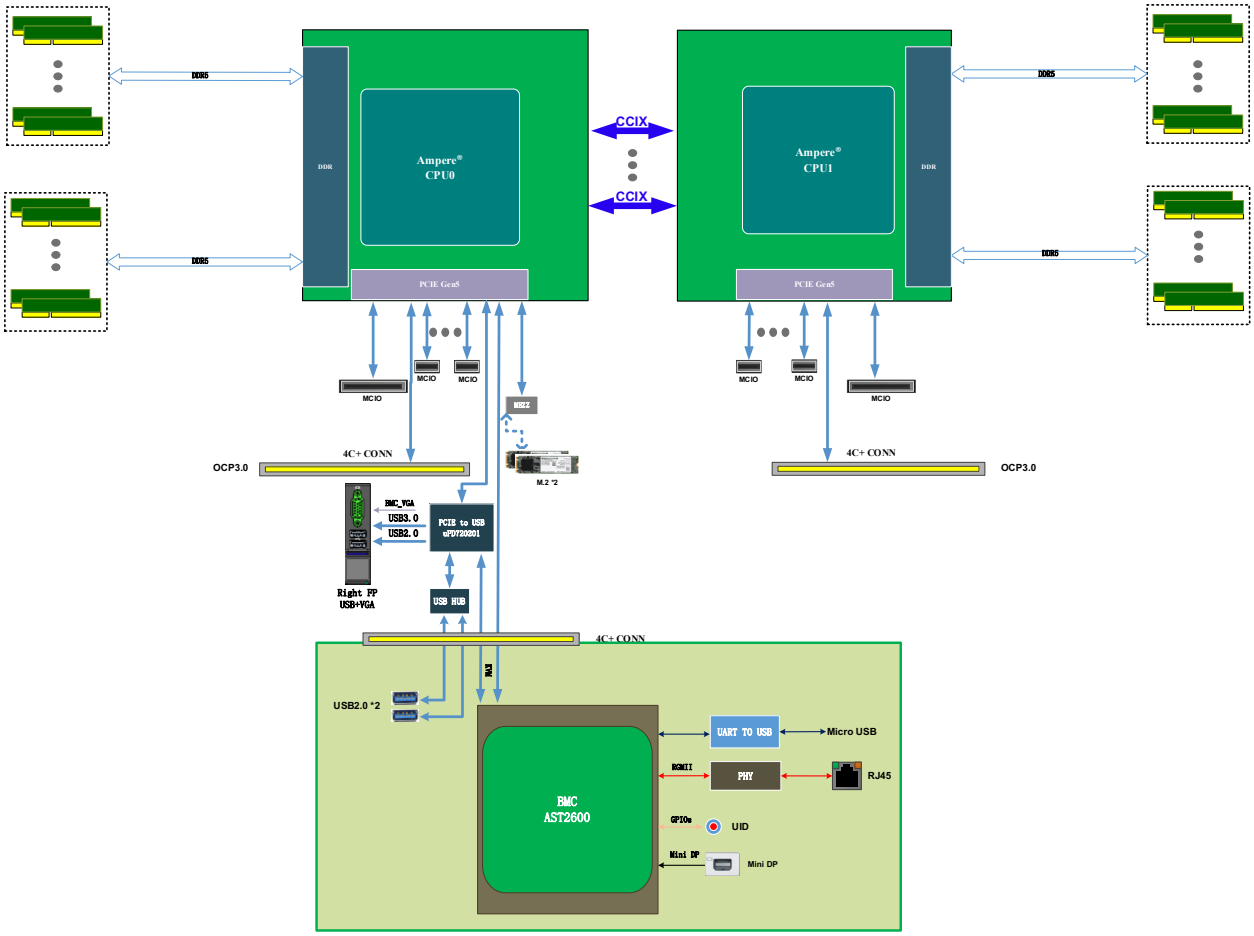

Figure 4: Mt. Mitchell Block Diagram

The server motherboard has the option to support two M.2 SSD drives. One can be used as boot drive and the other can be used as data drive or backup boot drive. Both the x4 PCIe link to the boot drive and the data drive shall be connected to the CPU.

Mt. Mitchell 2S server shall connect to an OCP v3.0 NIC on the platform through its PCIe. The server will utilize the remaining PCIe lanes to enable high bandwidth connections to the system board such as risers and backplane.

This document covers MB block diagram, features, BMC sub-system, power management, and miscellaneous physical specification of the motherboard.

6. Rack Compatibility

This motherboard will fit into a standard 19” rack chassis.

7. Physical Specifications

7.1. Motherboard Features

Mt. Mitchell motherboard supports for the following features:

Processor		
	CPU	Ampere® processor
	Sockets	Dual sockets or Single socket
Memory		
	DIMM Slots	16 or 24 or 32 DIMM slots total
	DIMM Type	DDR5 RDIMM with ECC
	DIMM Speed	DDR5
Storage		
	NVMe M.2	2x on-board 2280/22110 NVME SSDs, x4 Gen4
	NVMe U.2	2.5” Gen5 U.2 NVMe drives with hot plug support
	E3.S	2.5” Gen5 E3.S SSD drives with hot plug support
	SAS/SATA	3.5” SAS/SATA drives with hot plug support
	E1.S	Gen5 E1.S SSD drives with hot plug support at rear
PCI-Express Expansion		
	PCIe Risers	Support up to 6x FHFL PCIE cards and 2x FHHL PCIE cards
	MCIO Connectors	Gen5 connectors
	OCP v3.0	2x Gen5 OCP3.0 connector with single-host and multi-host support
Networking		
	LOM	OCP3.0 NIC PCIe Gen5
	MGMT	1x 1GbE from BMC to Rear RJ45 Connector (system management module)
MISC		
	Front Panel	PWR/UID buttons with LED System Status LED One USB 2.0 port One USB 3.0 port One VGA connector

	Rear Panel	System Management Module: Mini-DP, RJ-45, 2x USB 2.0, UID Button with LED, Power/Status LEDs, Micro USB for BMC UART
	FAN	6x 6056 dual-rotor fan connectors Hot plug support
	PSU	Dual 80 Plus Redundant PSUs, up to 2000W with PMBus 1.2, hot plug support
Server Management		
	DC-SCM 1.0	BMC ASPEED AST2600
System Firmware		
	BIOS	AMI Aptio® V, EDK-II
	BMC	AMI MegaRAC®, OpenBMC
	Security	TPM 2.0

7.2. Motherboard Placement

Mt. Mitchell is a spread-core motherboard design. The motherboard is 424.9mm wide x 422.9mm deep and will fit in an EIA-19" 2U rack mount chassis. There are one set of fan connectors which are placed in the top side of the board to support six 6056 dual rotor fans for 2U chassis. The master CPU socket is placed in the right side of the motherboard which is required to boot the system when a single CPU is populated.

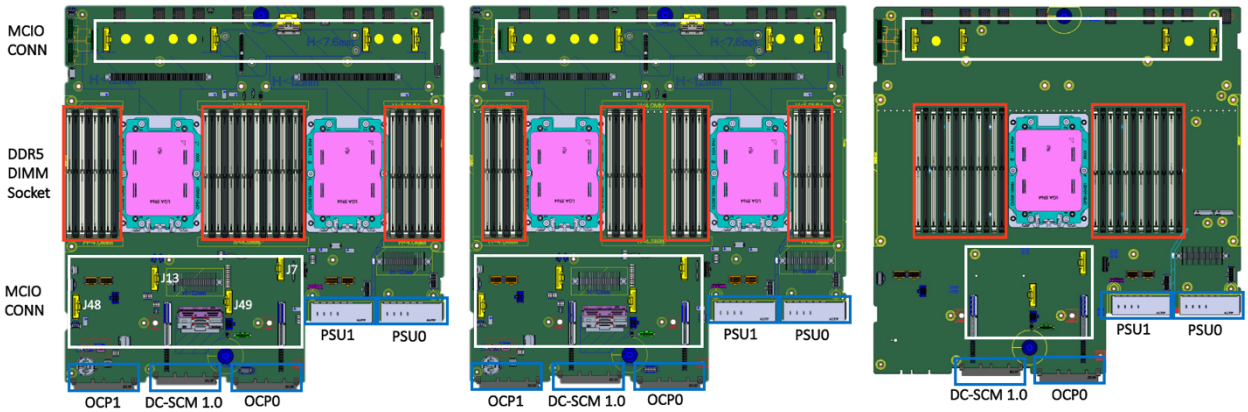

Figure 5: Mt. Mitchell Motherboard Placement

There are thirteen MCIO connectors which are placed on the motherboard. The MCIO connectors are used to connect CPU PCIe lanes to the NVMe backplane or OCP 3.0 card through cables. There are six connectors for CPU0 PCIe ports and the other seven connectors for CPU1 PCIe ports. One MCIO connector can be connected to OCP 3.0 card through cable to support x16 PCIE connection.

One M.2 on-board connector supports two 2280/22110 NVMe M.2 NVMe SSDs as system boot drives.

There are two MCIO connectors on Mt. Mitchell to support x16 PCIE risers.

There are two OCP 3.0 PCIE x16 connectors on Mt. Mitchell.

- For OCP0:
  - If J49 is not connected to other connectors, the status of OCP0 is a single-host x8
  - If J49 is connected to J7, the status of OCP0 is a single-host x16
  - If J49 is connected to J13, the status of OCP0 is a multi-host x16
- For OCP1:
  - If J48 is not connected to other connectors, the status of OCP1 is a single-host x8
  - If J48 is connected to J13, the status of OCP1 is a single-host x16
  - If J48 is connected to J7, the status of OCP1 is a multi-host x16

7.3. CPUs

Ampere® processor supports aarch64-compatible cores, DDR5 DIMM, PCIE Gen5, and miscellaneous processor interfaces for system management and control.

The detail PCIe configurations of Mt. Mitchell board are described in PCIe Configuration section.

Mt. Mitchell motherboard supports single or dual Ampere® processors.

7.4. DIMMs

Mt. Mitchell motherboard has three architecture options, the first one is two sockets, with total 8 DDR channels per socket, support up to 32 DIMMs. The second one is two sockets, with total 12 DDR channels per socket, support up to 24 DIMMs. The last one is one socket, with total 12 DDR channels, support up to 24 DIMM.
The Figure 6 shows the labels of 32 DIMM sockets which is used for DIMM installation order.

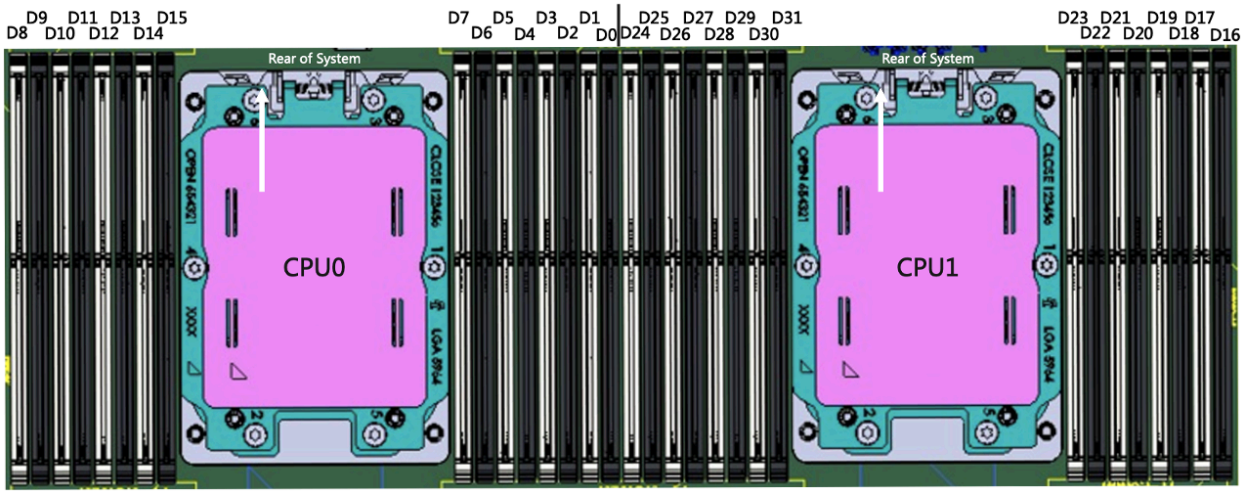

Figure 6 Mt. Mitchell DIMM Socket Label

The Table 1 & Table 2 show DIMM installation sequence for each socket, Master and Slave.

DIMM Configurations (Socket0)	DIMM Slots															
	D8	D9	D10	D11	D12	D13	D14	D15	D7	D6	D5	D4	D3	D2	D1	D0
2 DIMM															V	V
4 DIMM	V	V													V	V
8 DIMM	V	V	V	V									V	V	V	V
12 DIMM	V	V	V	V	V	V					V	V	V	V	V	V
16 DIMM	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V

Table 1 Installation Sequence for DIMMs in Socket0 (Master) side

DIMM Configurations (Socket1)	DIMM Slots															
	D24	D25	D26	D27	D28	D29	D30	D31	D23	D22	D21	D20	D19	D18	D17	D16
2 DIMM															V	V
4 DIMM	V	V													V	V
8 DIMM	V	V	V	V									V	V	V	V
12 DIMM	V	V	V	V	V	V					V	V	V	V	V	V
16 DIMM	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V

Table 2 Installation Sequence for DIMMs in Socket1 (Slave) side

The figure below shows the labels of 24 DIMM sockets which is used for DIMM installation order.

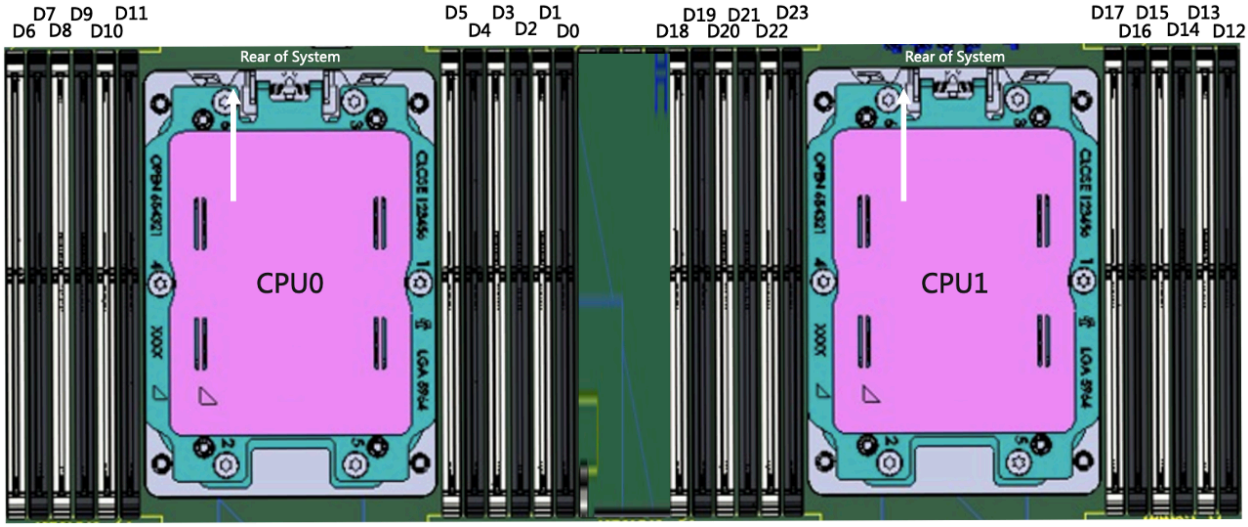

Figure 7 Mt. Mitchell DIMM Socket Label

The tables below show DIMM installation sequence for each socket, Master and Slave.

DIMM Configurations (Socket0)	DIMM Slots											
	D6	D7	D8	D9	D10	D11	D5	D4	D3	D2	D1	D0
2 DIMM											V	V
4 DIMM	V	V									V	V
8 DIMM	V	V	V	V					V	V	V	V
12 DIMM	V	V	V	V	V	V	V	V	V	V	V	V

Table 3 Installation Sequence for DIMMs in Socket0 (Master) side

DIMM Configurations (Socket1)	DIMM Slots											
	D18	D19	D20	D21	D22	D23	D17	D16	D15	D14	D13	D12
2 DIMM											V	V
4 DIMM	V	V									V	V
8 DIMM	V	V	V	V					V	V	V	V
12 DIMM	V	V	V	V	V	V	V	V	V	V	V	V

Table 4 Installation Sequence for DIMMs in Socket1 (Slave) side

The Figure below shows the labels of 24 DIMM sockets which is used for DIMM installation order.

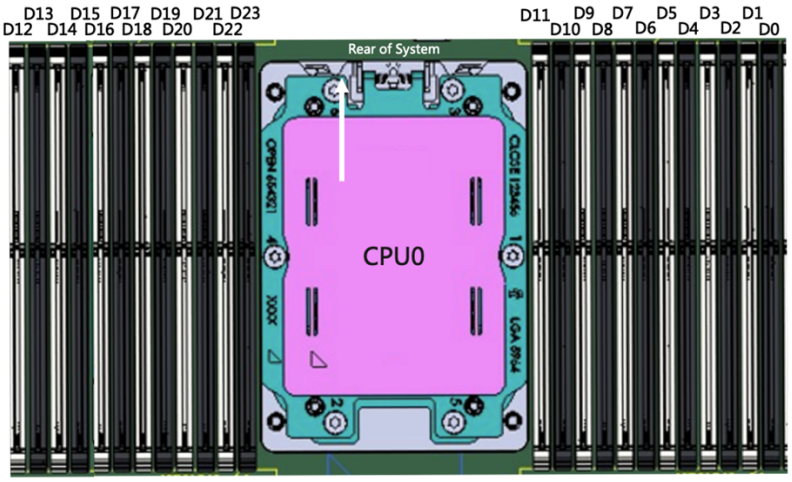

Figure 8 Mt. Mitchell DIMM Socket Label

The tables below show DIMM installation sequence for each socket, Master and Slave.

DIMM Configurations (Socket0)	DIMM Slots																					
	D12	D13	D14	D15	D16	D17	D18	D19	D20	D21	D22	D23	D11	D10	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0
2 DIMM																							V	V
4 DIMM	V	V																					V	V
8 DIMM	V	V	V	V																	V	V	V	V
12 DIMM	V	V	V	V	V	V													V	V	V	V	V	V
16 DIMM	V	V	V	V	V	V	V	V									V	V	V	V	V	V	V	V
20 DIMM	V	V	V	V	V	V	V	V	V	V					V	V	V	V	V	V	V	V	V	V
24 DIMM	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V

Table 5 Installation Sequence for DIMMs in Socket0 (Master) side

7.5. PCIe Configuration

Ampere® processor has Gen5 PCIe Root Complex.

7.5.1. Mt. Mitchell PCIE Configuration

The Table 6 shows the Mt. Mitchell PCIE ports configuration, device connectivity.

Mt. Mitchell PCIE Configuration		
CPU	**PCIE Port Connectors**	**PCIE Port Connectivity**
S0_RC0	No Connector	Inter processor link
S0_RC1	M.2 Connector MCIO Right Angle Connector	1 M.2 Connector for 2 M.2 SSD 1 MCIO
S0_RC2	MCIO Vertical Connector	2 MCIO
S0_RC3		2 MCIO
S0_RC4	BMC Connector	PCIE to USB, no connector PCIE to BMC, no connector
S0_RC5	No Connector	Inter processor link
S0_RC6	MCIO Vertical Connector OCP Connector	1 MCIO 1 OCPv3.0
S0_RC7	MCIO Vertical Connector	1 MCIO
S1_RC0	MCIO Vertical Connector	2 MCIO
S1_RC1		2 MCIO
S1_RC2		2 MCIO
S1_RC3	No Connector	Inter processor link
S1_RC4	MCIO Vertical Connector	1 MCIO
S1_RC5	MCIO Vertical Connector OCP Connector	1 MCIO 1 OCPv3.0
S1_RC6	No Connector	Inter processor link
S1_RC7	No Connector	No Connection

Table 6 Mt. Mitchell MB PCIE Ports Configuration

7.5.2. PCIE connectors

Mt. Mitchell MB uses the different MCIO connectors to support Gen5 PCIe interfaces and M.2 riser connector for dual M.2 NVMe SSD devices:

- Mini Cool Edge I/O vertical connector for NVME-U.2 SSD drives and OCP 3.0 PCIE connection
- Mini Cool Edge I/O right angle connector for NVME U.2 SSD drives
- Mini Cool Edge I/O x16 vertical connector for PCIe riser cards
- SFF-TA-1002 4C+ connector for OCP 3.0 NIC
- M.2 riser connector for M.2 NVMe SSD drives

7.5.2.1. MCIO PCIE x8 vertical Connector and Pin definition

Figure 9 AMPHENOL: G97V22312HR Pin:74

PIN	SIGNAL	PIN	SIGNAL
A1	GND	B1	GND
A2	P5E_RX_DP<7>	B2	P5E_TX_C_DP<7>

PIN	SIGNAL	PIN	SIGNAL
A3	P5E_RX_DN<7>	B3	P5E_TX_C_DN<7>
A4	GND	B4	GND
A5	P5E_RX_DP<6>	B5	P5E_TX_C_DP<6>
A6	P5E_RX_DN<6>	B6	P5E_TX_C_DN<6>
A7	GND	B7	GND
A8	MCIO_ID0	B8	ID0_ADDR1
A9	MCIO_WAKE0_N	B9	ID0_ADDR2
A10	GND	B10	ID0_ADDR3
A11	CLK_100M _REF0_DP	B11	MCIO_00_PERST_L
A12	CLK_100M_ REF0_DN	B12	ID0_ADDR0
A13	GND	B13	GND
A14	P5E_RX_DP<5>	B14	P5E_TX_C_DP<5>
A15	P5E_RX_DN<5>	B15	P5E_TX_C_DN<5>
A16	GND	B16	GND
A17	P5E_RX_DP<4>	B17	P5E_TX_C_DP<4>
A18	P5E_RX_DN<4>	B18	P5E_TX_C_DN<4>
A19	GND	B19	GND
A20	P5E_RX_DP<3>	B20	P5E_TX_C_DP<3>
A21	P5E_RX_DN<3>	B21	P5E_TX_C_DN<3>
A22	GND	B22	GND
A23	P5E_RX_DP<2>	B23	P5E_TX_C_DP<2>
A24	P5E_RX_DN<2>	B24	P5E_TX_C_DN<2>
A25	GND	B25	GND
A26	MCIO_ID1	B26	ID1_ADDR1
A27	MCIO_WAKE1_N	B27	ID1_ADDR2
A28	GND	B28	ID1_ADDR3
A29	CLK_100M_REF1_DP	B29	MCIO_01_PERST_L
A30	CLK_100M_REF1_DN	B30	ID1_ADDR0
A31	GND	B31	GND
A32	P5E_RX_DP<1>	B32	P5E_TX_C_DP<1>
A33	P5E_RX_DN<1>	B33	P5E_TX_C_DN<1>

PIN	SIGNAL	PIN	SIGNAL
A34	GND	B34	GND
A35	P5E_RX_DP<0>	B35	P5E_TX_C_DP<0>
A36	P5E_RX_DN<0>	B36	P5E_TX_C_DN<0>
A37	GND	B37	GND

Table 7 MCIO PCIE x8 Vertica/Right Angle Connector Pin Definition

7.5.2.2. MCIO PCIE x8 right angle Connector and Pin definition

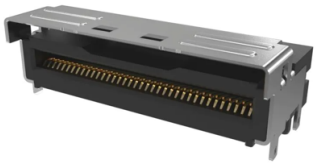

Figure 10 AMPHENOL: G97R22312HR Pin:74

7.5.2.3. MCIO PCIE x16 vertical Connector and Pin definition

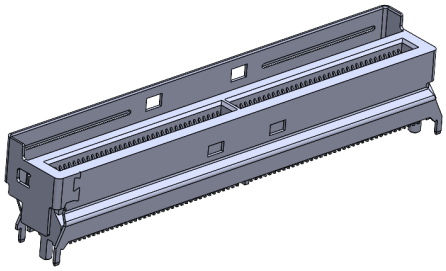

Figure 11 AMPHENOL: G97V26312HR Pin:148

PIN	SIGNAL	PIN	SIGNAL
A1	GND	B1	GND
A2	P5E_RX_DP<15>	B2	P5E_TX_C_DP<15>
A3	P5E_RX_DN<15>	B3	P5E_TX_C_DN<15>
A4	GND	B4	GND
A5	P5E_RX_DP<14>	B5	P5E_TX_C_DP<14>
A6	P5E_RX_DN<14>	B6	P5E_TX_C_DN<14>
A7	GND	B7	GND
A8	MCIO_ID0	B8	ID0_ADDR1
A9	MCIO_WAKE0_N	B9	ID0_ADDR2
A10	GND	B10	ID0_ADDR3

PIN	SIGNAL	PIN	SIGNAL
A11	CLK_100M _REF0_DP	B11	PERST0_L
A12	CLK_100M_ REF0_DN	B12	ID0_ADDR0
A13	GND	B13	GND
A14	P5E_RX_DP<13>	B14	P5E_TX_C_DP<13>
A15	P5E_RX_DN<13>	B15	P5E_TX_C_DN<13>
A16	GND	B16	GND
A17	P5E_RX_DP<12>	B17	P5E_TX_C_DP<12>
A18	P5E_RX_DN<12>	B18	P5E_TX_C_DN<12>
A19	GND	B19	GND
A20	P5E_RX_DP<11>	B20	P5E_TX_C_DP<11>
A21	P5E_RX_DN<11>	B21	P5E_TX_C_DN<11>
A22	GND	B22	GND
A23	P5E_RX_DP<10>	B23	P5E_TX_C_DP<10>
A24	P5E_RX_DN<10>	B24	P5E_TX_C_DN<10>
A25	GND	B25	GND
A26	MCIO_ID1	B26	ID1_ADDR1
A27	NC	B27	ID1_ADDR2
A28	GND	B28	ID1_ADDR3
A29	CLK_100M_REF1_DP	B29	PERST1_L
A30	CLK_100M_REF1_DN	B30	ID1_ADDR0
A31	GND	B31	GND
A32	P5E_RX_DP<9>	B32	P5E_TX_C_DP<9>
A33	P5E_RX_DN<9>	B33	P5E_TX_C_DN<9>
A34	GND	B34	GND
A35	P5E_RX_DP<8>	B35	P5E_TX_C_DP<8>
A36	P5E_RX_DN<8>	B36	P5E_TX_C_DN<8>
A37	GND	B37	GND
A38	GND	B38	GND
A39	P5E_RX_DP<7>	B39	P5E_TX_C_DP<7>
A40	P5E_RX_DN<7>	B40	P5E_TX_C_DN<7>
A41	GND	B41	GND

PIN	SIGNAL	PIN	SIGNAL
A42	P5E_RX_DP<6>	B42	P5E_TX_C_DP<6>
A43	P5E_RX_DN<6>	B43	P5E_TX_C_DN<6>
A44	GND	B44	GND
A45	NC	B45	ID0_ADDR1
A46	NC	B46	ID0_ADDR2
A47	GND	B47	ID0_ADDR3
A48	CLK_100M _REF2_DP	B48	PERST2_L
A49	CLK_100M_ REF2_DN	B49	ID0_ADDR0
A50	GND	B50	GND
A51	P5E_RX_DP<5>	B51	P5E_TX_C_DP<5>
A52	P5E_RX_DN<5>	B52	P5E_TX_C_DN<5>
A53	GND	B53	GND
A54	P5E_RX_DP<4>	B54	P5E_TX_C_DP<4>
A55	P5E_RX_DN<4>	B55	P5E_TX_C_DN<4>
A56	GND	B56	GND
A57	P5E_RX_DP<3>	B57	P5E_TX_C_DP<3>
A58	P5E_RX_DN<3>	B58	P5E_TX_C_DN<3>
A59	GND	B59	GND
A60	P5E_RX_DP<2>	B60	P5E_TX_C_DP<2>
A61	P5E_RX_DN<2>	B61	P5E_TX_C_DN<2>
A62	GND	B62	GND
A63	NC	B63	RISER_SCL
A64	THROTTLE_N	B64	RISER_SDA
A65	GND	B65	PRSNT_N
A66	CLK_100M _REF3_DP	B66	PERST3_L
A67	CLK_100M_ REF3_DN	B67	NC
A68	GND	B68	GND
A69	P5E_RX_DP<1>	B69	P5E_TX_C_DP<1>
A70	P5E_RX_DN<1>	B70	P5E_TX_C_DN<1>
A71	GND	B71	GND
A72	P5E_RX_DP<0>	B72	P5E_TX_C_DP<0>

PIN	SIGNAL	PIN	SIGNAL
A73	P5E_RX_DN<0>	B73	P5E_TX_C_DN<0>
A74	GND	B74	GND

Table 8 MCIO PCIE x16 Vertica Connector Pin Definition

7.5.2.4. SFF-TA-1002 4C+ connector

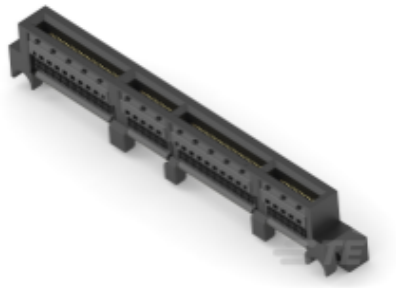

Figure 12 TE: 22340321-1 PIN: 168

PIN	SIGNAL	PIN	SIGNAL
OA1	PERST2_N	OB1	PWRGD
OA2	PERST3_N	OB2	PWR_EN
OA3	WAKE_N	OB3	OCP_LD
OA4	NCSI_OCP_ARBIN	OB4	NC
OA5	NCSI_OCP_ARBOUT	OB5	OCP_DATA_OUT
OA6	OCP_SLOT_ID1	OB6	OCP_CLK
OA7	OCP_TXEN	OB7	OCP_SLOT_ID0
OA8	OCP_TXD1	OB8	NCSI_OCP_RXD1
OA9	OCP_TXD0	OB9	NCSI_OCP_RXD0
OA10	GND	OB10	GND
OA11	CLK_100M_REF3_DN	OB11	CLK_100M_REF2_DN
OA12	CLK_100M_REF3_DP	OB12	CLK_100M_REF2_DP
OA13	GND	OB13	GND
OA14	CLK_50M_OCP	OB14	OCP_CRSDV
A1	GND	B1	P12V_STBY_OCP
A2	GND	B2	P12V_STBY_OCP
A3	GND	B3	P12V_STBY_OCP

PIN	SIGNAL	PIN	SIGNAL
A4	GND	B4	P12V_STBY_OCP
A5	GND	B5	P12V_STBY_OCP
A6	GND	B6	P12V_STBY_OCP
A7	SMB_SCL	B7	OCP_BIF0
A8	SMB_SDA	B8	OCP_BIF1
A9	RST_SMB_N	B9	OCP_BIF2
A10	PRSNT_A	B10	PERST0_N
A11	PERST1_N	B11	P3V3_STBY_OCP
A12	PRSNT_B2	B12	STBY_PWR_EN
A13	GND	B13	GND
A14	CLK_100M_REF1_DN	B14	CLK_100M_REF0_DN
A15	CLK_100M_REF1_DP	B15	CLK_100M_REF0_DP
A16	GND	B16	GND
A17	P5E_RX_DP<0>	B17	P5E_TX_C_DP<0>
A18	P5E_RX_DN<0>	B18	P5E_TX_C_DN<0>
A19	GND	B19	GND
A20	P5E_RX_DP<1>	B20	P5E_TX_C_DP<1>
A21	P5E_RX_DN<1>	B21	P5E_TX_C_DN<1>
A22	GND	B22	GND
A23	P5E_RX_DP<2>	B23	P5E_TX_C_DP<2>
A24	P5E_RX_DN<2>	B24	P5E_TX_C_DN<2>
A25	GND	B25	GND
A26	P5E_RX_DP<3>	B26	P5E_TX_C_DP<3>
A27	P5E_RX_DN<3>	B27	P5E_TX_C_DN<3>
A28	GND	B28	GND
A29	GND	B29	P5E_TX_C_DP<4>
A30	P5E_RX_DP<4>	B30	P5E_TX_C_DN<4>
A31	P5E_RX_DN<4>	B31	GND
A32	GND	B32	P5E_TX_C_DP<5>
A33	P5E_RX_DP<5>	B33	P5E_TX_C_DN<5>
A34	P5E_RX_DN<5>	B34	GND

PIN	SIGNAL	PIN	SIGNAL
A35	GND	B35	P5E_TX_C_DP<6>
A36	P5E_RX_DP<6>	B36	P5E_TX_C_DN<6>
A37	P5E_RX_DN<6>	B37	GND
A38	GND	B38	P5E_TX_C_DP<7>
A39	P5E_RX_DP<7>	B39	P5E_TX_C_DN<7>
A40	P5E_RX_DN<7>	B40	GND
A41	GND	B41	GND
A42	PRSNT_B1	B42	PRSNT_B0
A43	GND	B43	GND
A44	P5E_RX_DP<8>	B44	P5E_TX_C_DP<8>
A45	P5E_RX_DN<8>	B45	P5E_TX_C_DN<8>
A46	GND	B46	GND
A47	P5E_RX_DP<9>	B47	P5E_TX_C_DP<9>
A48	P5E_RX_DN<9>	B48	P5E_TX_C_DN<9>
A49	GND	B49	GND
A50	P5E_RX_DP<10>	B50	P5E_TX_C_DP<10>
A51	P5E_RX_DN<10>	B51	P5E_TX_C_DN<10>
A52	GND	B52	GND
A53	P5E_RX_DP<11>	B53	P5E_TX_C_DP<11>
A54	P5E_RX_DN<11>	B54	P5E_TX_C_DN<11>
A55	GND	B55	GND
A56	P5E_RX_DP<12>	B56	P5E_TX_C_DP<12>
A57	P5E_RX_DN<12>	B57	P5E_TX_C_DN<12>
A58	GND	B58	GND
A59	P5E_RX_DP<13>	B59	P5E_TX_C_DP<13>
A60	P5E_RX_DN<13>	B60	P5E_TX_C_DN<13>
A61	GND	B61	GND
A62	P5E_RX_DP<14>	B62	P5E_TX_C_DP<14>
A63	P5E_RX_DN<14>	B63	P5E_TX_C_DN<14>
A64	GND	B64	GND
A65	P5E_RX_DP<15>	B65	P5E_TX_C_DP<15>

PIN	SIGNAL	PIN	SIGNAL
A66	P5E_RX_DN<15>	B66	P5E_TX_C_DN<15>
A67	GND	B67	GND
A68	NC	B68	OCP_LINK_LED
A69	NC	B69	NC
A70	PWRBRK	B70	PRSNT_B3

Table 9 OCP v3.0 Connector Pin Definition

7.5.2.5. M.2 Riser connector

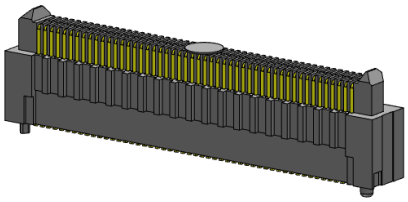

Figure 13 SAMTEC: ERM8-050-08.0-L-DV-K-TR PIN: 100

PIN	SIGNAL	PIN	SIGNAL
1	GND	2	GND
3	P5E_RX_DP<0>	4	P5E_TX_C_DP<0>
5	P5E_RX_DN<0>	6	P5E_TX_C_DN<0>
7	GND	8	GND
9	P5E_RX_DP<1>	10	P5E_TX_C_DP<1>
11	P5E_RX_DN<1>	12	P5E_TX_C_DN<1>
13	GND	14	GND
15	P5E_RX_DP<2>	16	P5E_TX_C_DP<2>
17	P5E_RX_DN<2>	18	P5E_TX_C_DN<2>
19	GND	20	GND
21	P5E_RX_DP<3>	22	P5E_TX_C_DP<3>
23	P5E_RX_DN<3>	24	P5E_TX_C_DN<3>
25	GND	26	GND
27	P5E_RX_DP<4>	28	P5E_TX_C_DP<4>
29	P5E_RX_DN<4>	30	P5E_TX_C_DN<4>

PIN	SIGNAL	PIN	SIGNAL
31	GND	32	GND
33	P5E_RX_DP<5>	34	P5E_TX_C_DP<5>
35	P5E_RX_DN<5>	36	P5E_TX_C_DN<5>
37	GND	38	GND
39	P5E_RX_DP<6>	40	P5E_TX_C_DP<6>
41	P5E_RX_DN<6>	42	P5E_TX_C_DN<6>
43	GND	44	GND
45	P5E_RX_DP<7>	46	P5E_TX_C_DP<7>
47	P5E_RX_DN<7>	48	P5E_TX_C_DN<7>
49	GND	50	GND
51	CLK_100M_REF0_DN	52	CLK_100M_REF1_DN
53	CLK_100M_REF0_DP	54	CLK_100M_REF1_DP
55	GND	56	GND
57	PERST_0_N	58	SCL
59	PERST_1_N	60	SDA
61	PRSNT_1_N	62	SALT10_N
63	PRSNT_1_N	64	WAKE_N
65	GND	66	GND
67	P3V3_M2	68	P3V3_M2
69	P3V3_M2	70	P3V3_M2
71	P3V3_M2	72	P3V3_M2
73	P3V3_M2	74	P3V3_STBY
75	P3V3_M2	76	P1V8_STBY
77	NC	78	NC
79	GND	80	GND
81	GND	82	GND
83	GND	84	GND
85	GND	86	GND
87	GND	88	GND
89	NC	90	NC
91	NC	92	NC

PIN	SIGNAL	PIN	SIGNAL
93	NC	94	NC
95	NC	96	NC
97	NC	98	NC
99	NC	100	NC

Table 10 M.2 Riser Connector Pin Definition

7.5.3. Mt. Mitchell PCIE Riser Cards

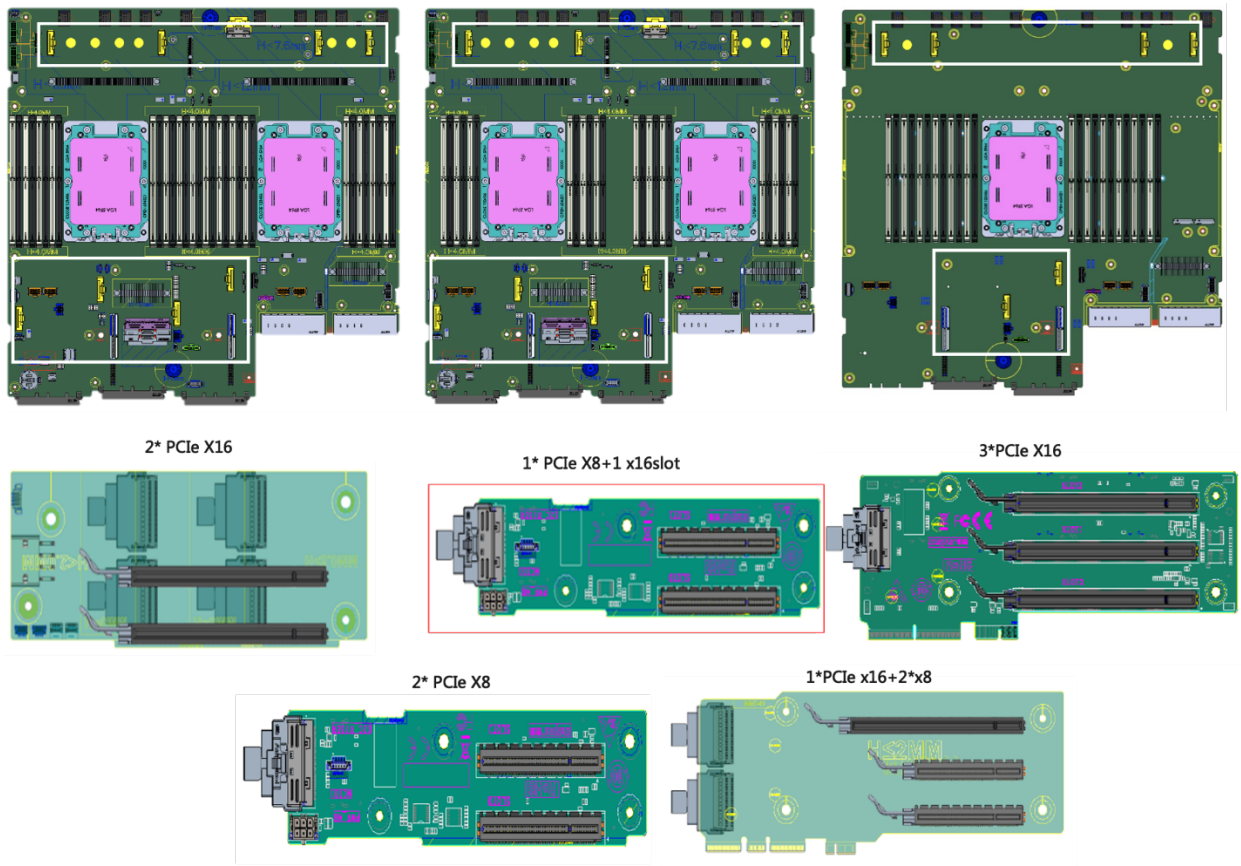

Figure 14 Mt. Mitchell 2U Riser Cards

7.6. Management Subsystem

7.6.1. BMC

The primary functions of BMC on Mt. Mitchell server board management subsystem are:

- IPMI2.0 support for remote management
- IPMB support for chassis management
- System power control
- System Event Log
- Sensor monitoring (voltage/power, temperature, PSU)
- Fan control (fail detection, speed monitor & control)
- Serial Over Lan
- Asset Information (FRU)
- Auto recovery from hang (Watchdog)
- Remote KVM and Media

Mt. Mitchell BMC is designed with AST2600 of which system is configured with 1024MB DDR4-3200 SDRAM, two 64MB SPI flash memories for BMC boot, 32GB eMMC storage, x1 Gen2 PCIE port for VGA, 14 I2C ports, 4 I3C ports and GPIO pins for system management, 1GbE management ETH port for network, UARTs for console and debug, USB for KVM/Redfish support, PWM/TACH for FAN control, and ADC for power fail monitoring. The Table 11 shows BMC configuration.

INTERFACES	BMC CONFIGURATION
Memory	1024MB DDR4 memory with ECC support
PCIe	VGA or DisplayPort support x1 PCIE Gen2, configured as EP, direct connection to Mt.Mitchell CRB socket0 PCIE
VGA & DisplayPort	1 VGA port with 1920x1028 resolution DisplayPort 1.1a output
SPI	2x 512Mb SPI flash for BMC boot, one for main and the other for fail-save,
Storage	1x 32GB eMMC device (optional use)
I2C	I2C_1 (Master): Board ID, Clock and Hot-Swap Controller I2C_[3:2] (Master): PSU, Power VRDs I2C_5 (Master): Temperature sensors, FRU EEPROM and ADC Controller on Mitchell CRB I2C_4 (Master): As a master/slave I2C, BMC requests both sockets' I2C slave for SOC side band I2C_[9:6] (Master): Fan Controller, PCIe Gen-Z, Front Panel, RTC I2C_10 (Master): NVMe hot-plug, Backplane, IO expander for PCIe present and power break functions I2C_12 (Slave): BMC receives Socket0's IPMI SSIF requests
Network	x1 10/100/1000Mbps management LAN from BMC's ETH x1 10/100 RMII interface as NCSI for OCPv3 NIC

INTERFACES	BMC CONFIGURATION
UART	One 4-pin UART port Six 2-pin UART ports
USB	One USB port, directly connected to PCIE-USB bridge chip One USB port connects to HDR1x4
GPIO	Refer *Table 12* for GPIO pin assignments.
ADC	Monitor power rails of motherboard
PWM/TACH	5 fans for 2U systems
JTAG	Support one JTAG master port. This port will be used to program CPLD, FPGA (PCIe FPGA card). It can be used to debug two Mt.Mitchell CRB Socket if debug software is available on BMC

Table 11 BMC AST2600 Configuration

7.6.2. I2C Network

Mt. Mitchell MB has many I2C devices onboard such as voltage regulator devices, EEPROMs (DDR SPD, FRU), Temperature sensors, RTC, PSU, and PCIe backplane controllers & PCIe I2C side band signals. Also, Ampere® processor uses I2C interfaces to communicate with BMC.

Mt. Mitchell has 20 I2C interfaces with speeds of up to 1 MHz. An interface can be a master or slave (statically).

- I2C9: Side Communication channel
- I2C6: RTC
- I2C5: PCIe Hot Plug
- I2C4, 7: Connect to other slave devices or headers
- I2C_BMC: External BMC – SECpro/Mpro receive BMC requests for SOC sideband
- I2C2: External BMC – SECpro/Mpro respond to BMC requests for IPMI/SSIF
- I2C1: Reserved
- I2C0: Reserved
- I2C_VRDEXT: PCP VRD
- I2C3: Other VRDs

Mt. Mitchell employs 14x I2C buses and 4x I3C buses of BMC to manage the board:

- I2C_1 (Master): Board ID, Clock and Hot-Swap Controller
- I2C_2 (Master): Power VRDs
- I2C_3 (Master): PSU, some main power rails, and efuses on Mt. Mitchell Mainboard
- I2C_5 (Master): Temperature sensors, FRU EEPROM, TPM, Front panel LEDs and ADC Controller on Mt. Mitchell Mainboard

- I2C_4 (Master/Slave): As a master/slave I2C, BMC requests both sockets' I2C slave for SOC side band
- I2C_6 (Master): USB Hub, PCIe OCP and Riser cards.
- I2C_7 (Master): RTC
- I2C_8 (Master): Front Panel
- I2C_9 (Master): Fan Controller and IO expander for Fan present
- I2C_10 (Master): NVMe hot-plug, Front Backplane, Rear Backplane, IO expander for power throttle functions
- I2C_12 (Slave): BMC receives CPU0's IPMI SSIF requests
- I2C_15 (Master): Temperature sensors and FRU ID EEPROM on DC-SCM
- I2C_16 (Master): Connected to CPLD on Mt. Mitchell
- I3C_1/2/3/4 (Master): DDR SPD

7.6.2.2. I2C connection of DIMM SPDs

Eight I2C controllers, I2C_DDR_[7:0], in Ampere® processor are used to access the DIMM SPD EEPROMs. Each DDR channel of which SPD EEPROMs are attached to a dedicated I2C controller, which allows to use same I2C addresses for the SPD EEPROM on each DIMM channel.

Four I3C controllers, I3C_[4:1], in BMC are used to access the DIMM SPD EEPROM as well. Four DDR channels form a group of which SPD EEPROMs are attached to a dedicated I3C controller, which allows to use same I3C addresses for the SPD EEPROM on each DIMM group.

7.6.2.3. I2C connection between CPU and BMC

CPU and BMC communicates via asymmetric I2C connectivity. When CPU sends requests to BMC, the requests will be sent from the I2C_2 of the CPU in socket 0 as master port to the I2C_12 of the BMC as slave port. However, when BMC sends requests to CPU, the requests will be sent from the I2C_4 of the BMC as master port to the I2C_BMC slave port, either the one in Socket 0 or the other in Socket 1. BMC may probe Mpro side-band (MCTP or PLDM) from I2C_4.

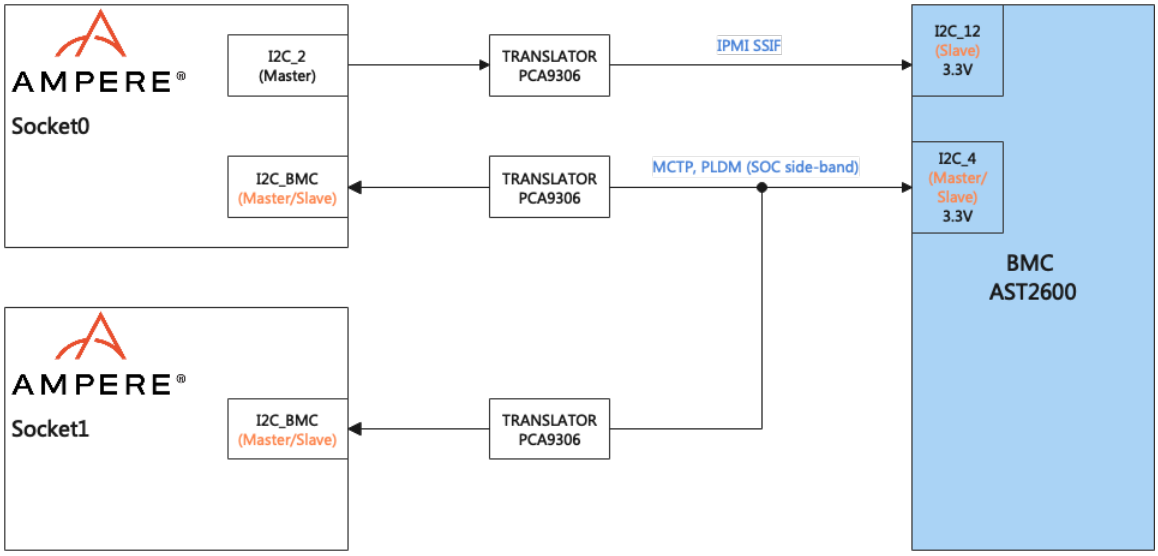

Figure 15 Mt. Mitchell I2C connection between CPU and BMC

7.6.2.4. I2C connections for PCIe NVMe hot plug & OOB management, PCIE sideband signals

Figure 16 shows I2C connections for PCIE NVMe hot-plug and sideband signals for PCIE devices. The I2C_5 interface of Ampere® processor is connected to I/O expander on NVMe backplane to support hot-plug feature. And BMC I2C_6 and I2C_10 interfaces are connected to PCIe sideband signals to support Wake signals and NVMe OOB management.

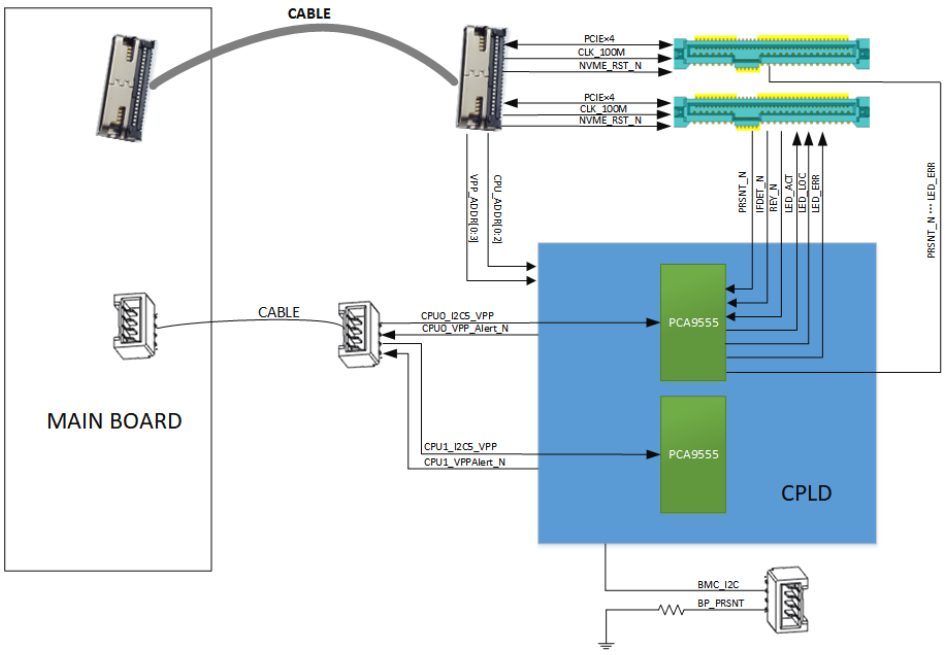

Figure 16 Mt. Mitchell I2C connections for PCIE NVMe hot-plug and management

7.6.2.5. Miscellaneous I2C devices

Temperature sensors, FRU, and PSU devices are connected to BMC. The RTC device is connected to Ampere® processor and BMC.

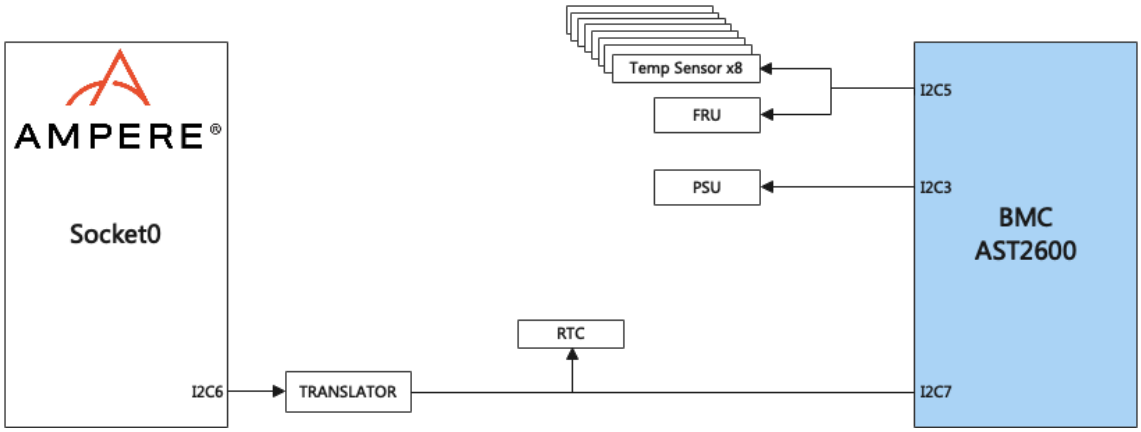

Figure 17 Mt. Mitchell Miscellaneous I2C devices such as temp-sensors, FRU, PSU, and RTC

7.7. Debug Interfaces

Mt. Mitchell MB has one dedicated JTAG header:

- Ampere® processor provides a single JTAG interface that can be daisy chained on 2P systems.
- BMC supports a JTAG master controller.
- Any JTAG Slave can be programmed through this JTAG port.
- The board must contain mux circuitry controlled by the BMC to switch the BMC to either CPLD or CPU's JTAGs which are the two programming paths.
- Note that, there must still be a CPLD JTAG Header for a backup plan to use external CPLD programming tool to program CPLD

By default, the CPLD JTAG slave port is connected to 1x8 header for CPLD Programmer, the BMC JTAG port is configured as slave and connected to 2x10 header for BMC JTAG debugger. BMC FW can configure the JTAG port as master and set GPIOM2 pin to connect the BMC JTAG port to CPLD JTAG slave port for updating CPLD image.

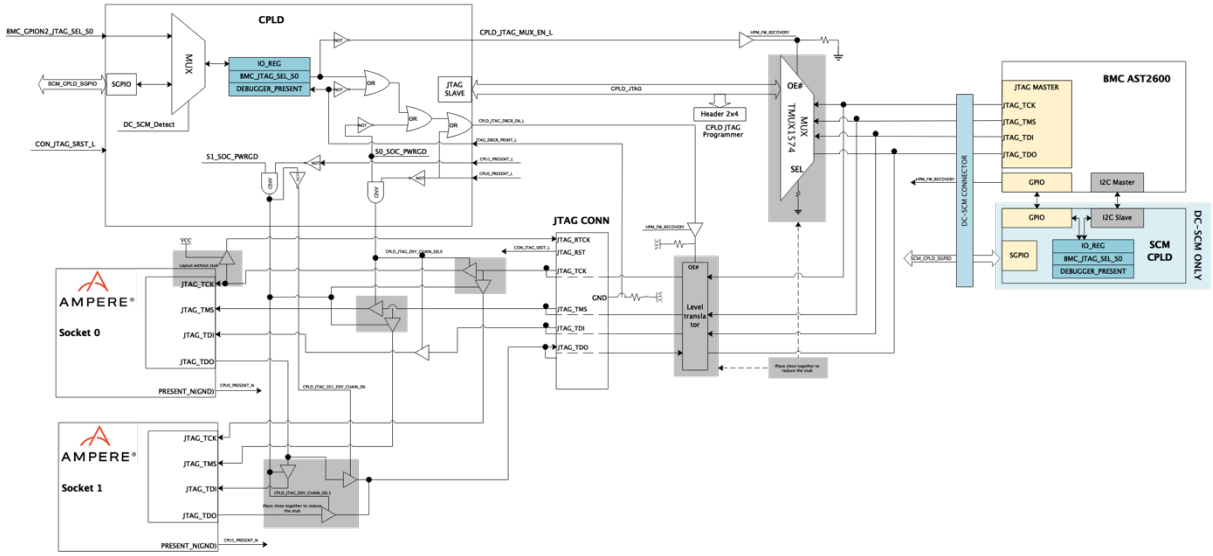

Figure 18 Mt. Mitchell JTAG Connectivity among CPU/CPLD/BMC

7.8. GPIOs

Figure 19 shows the connections between CPU and BMC GPIO pins. The CPLD provides simple voltage isolation since they're in different power rails. The detail definition of GPIO pins can be found in Table 12 and Table 15.

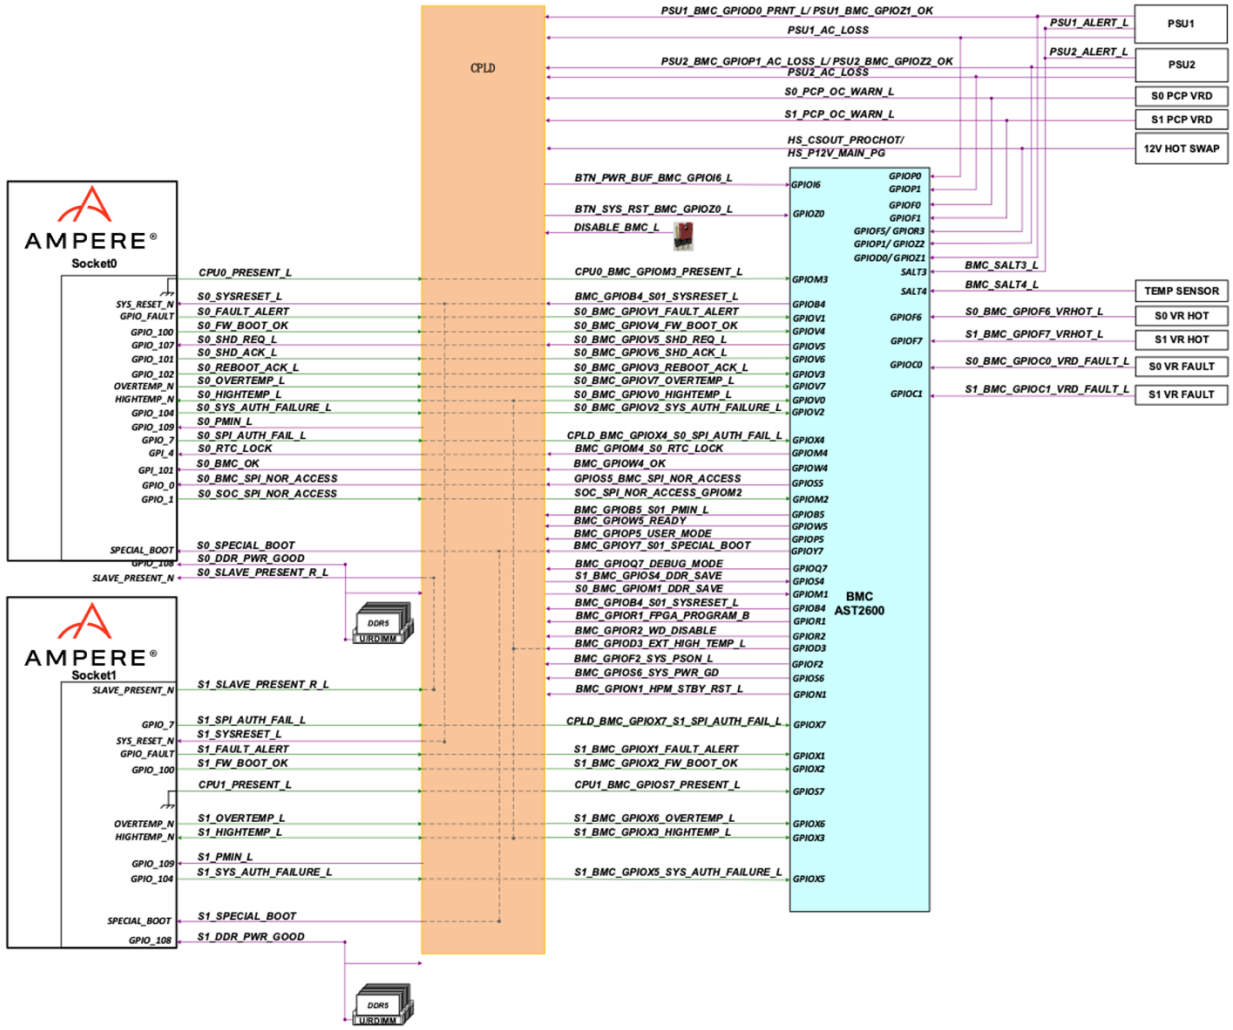

Figure 19 Mt. Mitchell GPIO mapping between CPU and BMC

CPU I/O		Socket 0	
PIN NAME	I/O (1.8V)	SIGNAL NAME	DESCRIPTION
GPIO_0	I, PD 10K	S0_BMC_SPI_NOR_ACCESS	Active High, connect to BMC via CPLD. This input indicates that the BMC is accessing the boot SPI NOR flash. Note that this PIN acts as grant to access SPI-NOR from the host point of view. The host does NOT access the SPI-NOR when this PIN is asserted.

GPIO_1	O, PD 10K	S0_SOC_SPI_NOR_ACCESS	Active High, connect to BMC via CPLD. This output indicates that SoC is accessing the boot SPI NOR flash. The BMC must yield to SoC until the pin is de-asserted. Note that this PIN acts as a request to access SPI-NOR from host point view. If the GPIO is not de-asserted within 500ms, host SPI-NOR operations fail.
GPIO_2	I/O, PD 1K	S0_GPIO2	Spare
GPIO_3	O, PU 4.7K	S0_GPIO3_CPLD_IOEXP_RST_L	Connect to CPLD for reset the PERST IO Expander
GPIO_4	I, PD 10K	S0_BIOS_RECOVER	S0 BIOS Recovery - Reserved for future use
GPIO_5	I/O, FLOATING	S0_GPIO5_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_6	I/O, FLOATING	S0_GPIO6_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_7	O, PU 4.7K	S0_SPI_AUTH_FAIL_L	S0 SPI boot authentication failure, active LOW. Connect to BMC via CPLD
GPIO_8	I/O, FLOATING	S0_GPIO8_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_9	O, PD 10K	S0_GPIO9_SPI_SOCSLV_SEL	Select S1 connect to SPI boot flash
GPIO_[15:10]	I/O, PD 10K	S0_GPI0[15:10]	Spare I/O
GPIO_[21:16]	I/O, FLOATING	S0_GPIO[21:16]_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_22	I, PU 4.7K	S0_PCIE_PRSNT_L	S0 PCIE card present. Interrupt from IO Expander, active LOW
GPIO_23	I/O, FLOATING	S0_GPI023_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_100	O, PD 10K	S0_FW_BOOT_OK	S0 SECpro boot OK, active HIGH. Connect to BMC via CPLD
GPIO_101	O, PU 4.7K	S0_SHD_ACK_L	S0 acknowledge “graceful shutdown", active LOW. Connect to BMC via CPLD
GPIO_102	O, PU 4.7K	S0_REBOOT_ACK_L	S0 software reboot acknowledge, active LOW. Connect to BMC via CPLD
GPIO_103	O, PD 10K	S0_DDR_SAVE_SPARE	S0 memory SAVE mode control by CPU, active HIGH. Connect to DIMM and BMC - Reserved for future use
GPIO_104	O, PU 4.7K	S0_SYS_AUTH_FAILURE_L	S0 SecureBoot authentication failure, active LOW. Connect to BMC via CPLD
GPIO_105	I, PD 10K	CPLD_S0_DDR_ADR_SPARE	S0 memory ADR, active HIGH - Reserved for future use
GPIO_106	I, PU 4.7K	S0_TPM_ALERT_L	TPM SPI1 module alert, active LOW
GPIO_107	I, PU 4.7K	S0_SHD_REQ_L	BMC request S0 a “graceful shutdown”, active LOW. Connect to BMC via CPLD
GPIO_108	I, PU 4.7K	S0_DDR_PWR_GOOD	S0 connect to DIMM PWR_GOOD module via CPLD
GPIO_109	I, PU 4.7K	S0_PMIN_L	S0 power limit, active LOW
GPI_0	I, PU	S0_GPI0_USER_MODE	Active high - Connect to CPLD - default HIGH
GPI_1	I, PU	S0_GPI1_DEBUG_MODE	Active low - Connect to CPLD - default HIGH
GPI_2	I, PU	S0_GPI2_WD_DISABLE	Active low - if set, firmware will disable watchdog to allow for uninterrupted debugging
GPI_3	I, PD 10K	S0_GPI3_SPARE_B	Spare

GPI_4	I, PD 10K	S0_RTC_LOCK	This signal indicates to S0 that the RTC access will be temporary restricted. BMC output HIGH this signal via CPLD to indicate that it needs access to the RTC - Active HIGH
GPI_[7:5]	I, PD 1K	S0_GPI[7:5]	Spare
GPI_100	I, PU 4.7K	S0_TPM_PRSNT_L	TPM SPI1 module present, active LOW
GPI_101	I, PD 10K	S0_BMC_OK	BMC is ready for operation, active HIGH. Connect to BMC via CPLD
SPECIAL_BOOT	I, PD 10K	S0_SPECIAL_BOOT	Jump to special boot mode. Control by BMC via CPLD
HIGHTEMP_N	O, PU 4.7K	S0_HIGHTEMP_L	S0 High temperature alert, active LOW. Connect to BMC via CPLD
OVERTEMP_N	O, PU 4.7K	S0_OVERTEMP_L	S0 Over temperature alert, BMC shutdown power immediately, active LOW. Connect to BMC via CPLD
GPIO_FAULT	O, PU 4.7K	S0_FAULT_ALERT_L	S0 Fault alert, control RED LED blinking, active LOW. Connect to BMC via CPLD

Table 12 GPIO Assignment Table for Master Socket

CPU I/O		Socket 1	
PIN NAME	I/O (1.8V)	PIN NAME	I/O (1.8V)
GPIO_[2:0]	I/O, PD 10K	S1_GPIO[2:0]	Spare I/O
GPIO_3	O, PU 4.7K	S1_GPIO3_CPLD_IOEXP_RST_L	Connect to CPLD for reset the PERST IO Expander
GPIO_4	I/O, PD 1K	S1_GPIO4	Spare I/O
GPIO_5	I/O, FLOATING	S1_GPIO5_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_6	I/O, PD 1K	S1_GPIO6	Spare I/O
GPIO_7	O, PU 4.7K	S1_SPI_AUTH_FAIL_L	S0 SPI boot authentication failure, active LOW. Connect to BMC via CPLD
GPIO_[11:8]	I/O, PD 1K	S1_GPIO[11:8]	Spare I/O
GPIO_[15:12]	O, PU 4.7K	S1_GPIO[15:2]	Spare I/O
GPIO_[21:16]	I/O, FLOATING	S1_GPIO[21:16]_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_22	I, PU 4.7K	S1_PCIE_PRSNT_L	S0 PCIE card present. Interrupt from IO Expander, active LOW
GPIO_23	I/O, FLOATING	S1_GPI023_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPIO_100	O, PD 10K	S1_FW_BOOT_OK	S0 SECpro boot OK, active HIGH. Connect to BMC via CPLD
GPIO_[102:101]	I/O, PD 1K	S1_GPIO[102:101]	Spare I/O
GPIO_103	O, PD 10K	S1_DDR_SAVE_SPARE	S1 memory SAVE mode control by CPU, active HIGH. Connect to DIMM and BMC - Reserved for future use
GPIO_104	O, PU 4.7K	S1_SYS_AUTH_FAILURE_L	S1 SecureBoot authentication failure, active LOW. Connect to BMC via CPLD

GPIO_105	I, PD 10K	CPLD_S1_DDR_ADR_SPARE	S1 memory ADR, active HIGH - Reserved for future use
GPIO_[107:106]	I/O, PD 1K	S1_GPIO[107:106]	Spare I/O
GPIO_108	I, PU 4.7K	S1_DDR_PWR_GOOD	S1 connect to DIMM PWR_GOOD module via CPLD
GPIO_109	I, PU 4.7K	S1_PMIN_L	S1 power limit, active LOW
GPI_[1:0]	I, PD 1K	S1_GPI[1:0]	Spare I/O
GPI_2	I, FLOATING	S1_GPI2_CPLD_SPARE	Connect to CPLD - Reserved for future use
GPI_3	I, PD 1K	S1_GPI3_SPARE_B	Spare
GPI_[7:4]	I, PD 1K	S1_GPI[7:4]	Spare
GPI_[101:100]	I, PD 1K	S1_GPIO[101:100]	Spare
SPECIAL_BOOT	I, PD 10K	S1_SPECIAL_BOOT	Jump to special boot mode. Control by BMC via CPLD
HIGHTEMP_N	O, PU 4.7K	S1_HIGHTEMP_L	S1 High temperature alert, active LOW. Connect to BMC via CPLD
OVERTEMP_N	O, PU 4.7K	S1_OVERTEMP_L	S1 Over temperature alert, BMC shutdown power immediately, active LOW. Connect to BMC via CPLD
GPIO_FAULT	O, PU 4.7K	S1_FAULT_ALERT_L	S1 Fault alert, control RED LED blinking, active LOW. Connect to BMC via CPLD

Table 13 GPIO Assignment Table for Slave Socket

7.9. Boot Devices

The below three devices are used for the system boot:

- Macronix MX25U51245GMI00 512Mbit QSPI flash for CPU SECpro/Mpro/ATF/UEFI FW
- Macronix MX25L51245GMI-08G 512Mbit QSPI flash for BMC AST2600 FW

The Figure 20 shows the boot devices connectivity between Ampere® processors and BMC. Two single bit SPI flash memories are attached to BMC, one for main boot device and the other for failover one.

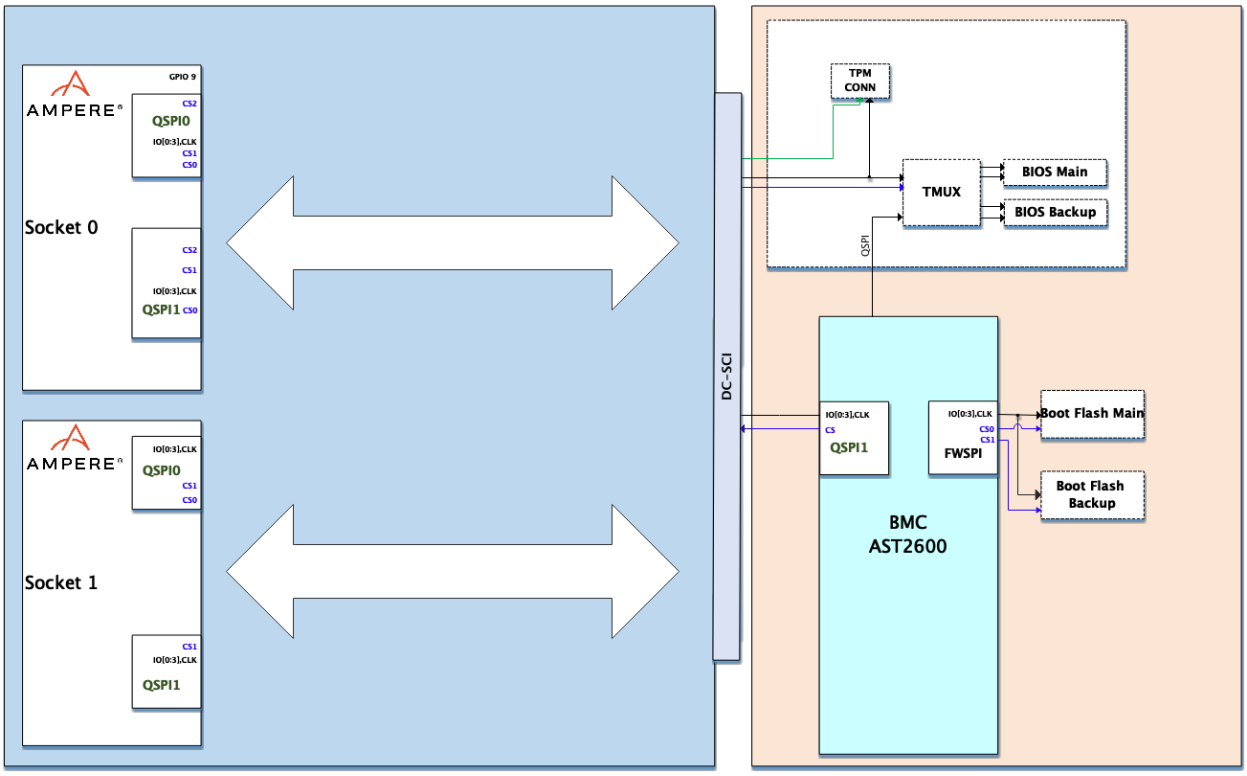

Figure 20 SPI Flash, and TPM for CPU and BMC

Two QSPI flash memories, main-boot and failover devices, are shared between Ampere® processor in master socket and BMC. While Ampere® processor fetches codes from flash memory in QSPI mode, BMC updates FW to flash memory in QSPI mode as well. BMC switches the QSPI flash CS signal from main-boot device to failover one though MUX controlled by GPIO pin. Also, BMC uses GPIO and MUX to gain access to the QSPI flash memory for FW update.

7.10. Clocks, Reset, and Power-on Sequence

7.10.1. Clocks

Figure 21 shows clock sources on Mt. Mitchell motherboard.

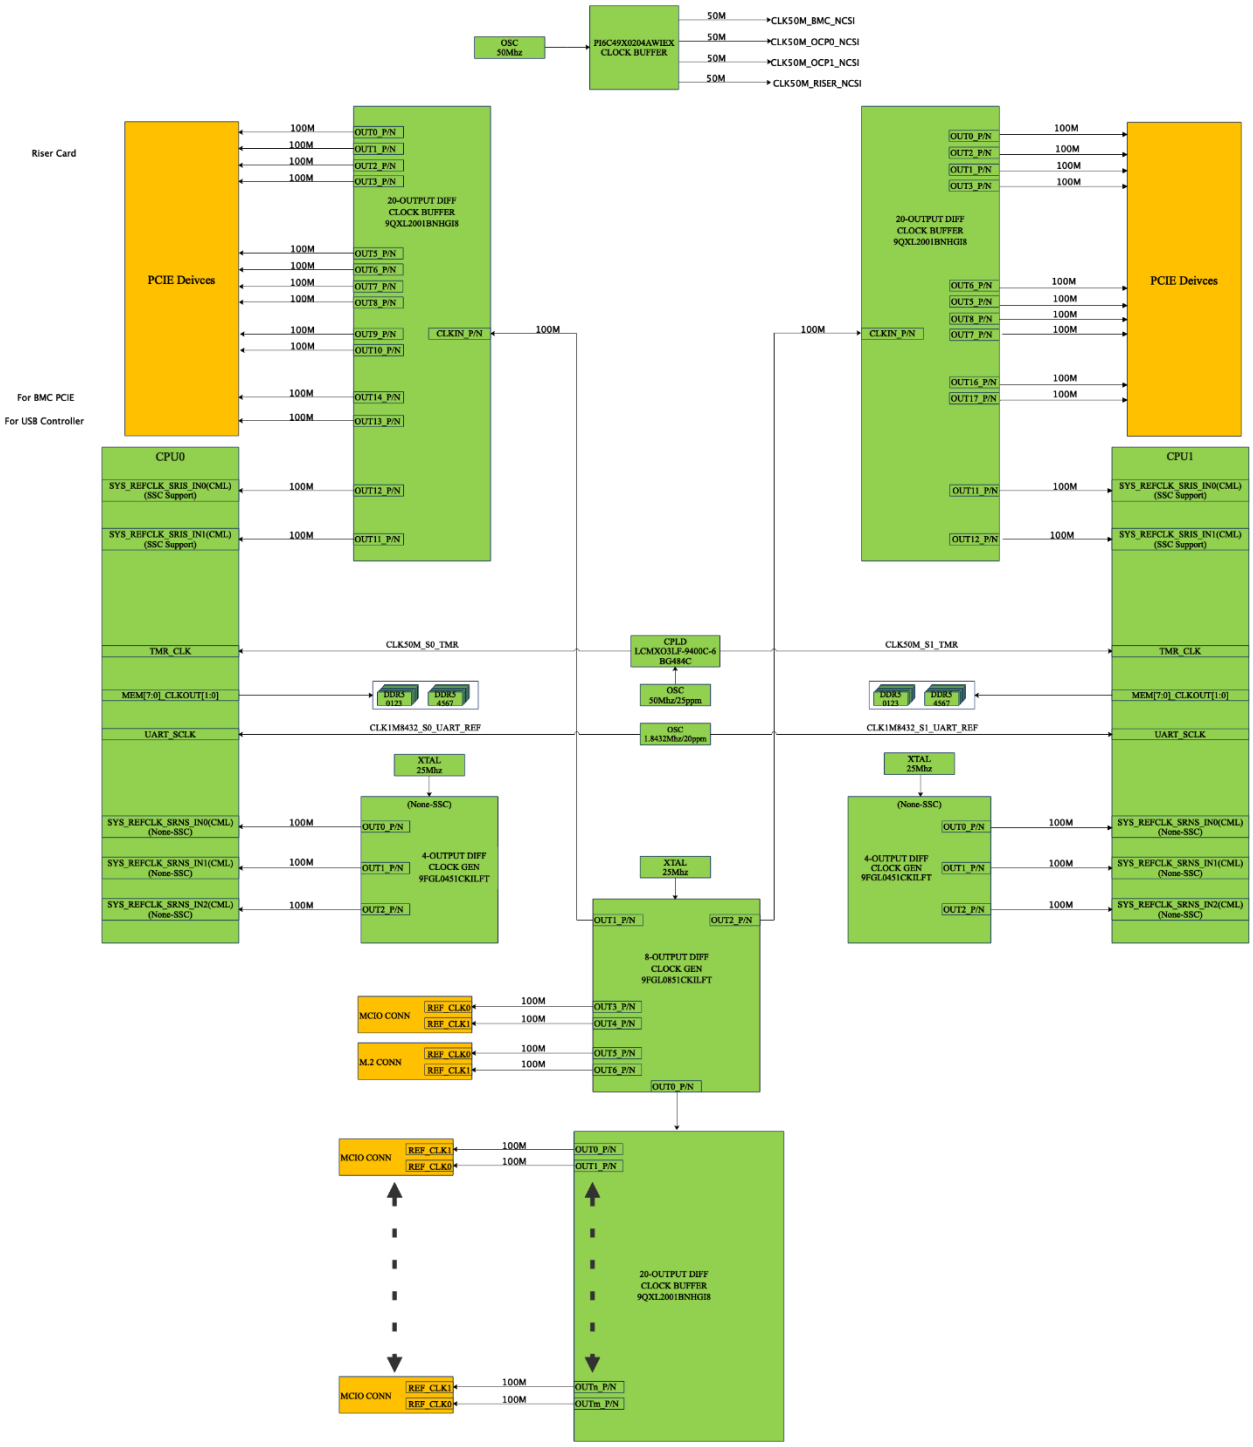

Figure 21 Mt. Mitchell Clock Sources

Renesas (IDT) 9FGL0451CKILFT 4-output clock generator chip provides Ampere® processor 100MHz PCIE Gen5 compliant clock without SSC as system reference clock and serdes reference clock for CCIX 2P link.

Renesas (IDT) 9FGL0851CKILFT 8-output clock generator and 9QXL2001BNHGI8 20-output clock buffer chip provide 100MHz PCIE Gen5 compliant clock for Ampere® processor, onboard PCIE devices, PCIE riser cards, OCPv3 NIC connector, and MCIO connectors for NVMe/E1.S drives as common reference clock.

CPLD takes 50MHz clock from oscillator and generates two 50MHz synchronous clocks to feed the processors in Master/Slave sockets as global timer clock.

Two 1.8432MHz oscillators are connected to Ampere® processors as UART serial clock sources to set UART baud rate to 115200 bps.

25MHz crystals are used as input clock sources for Renesas (IDT) Clock Generators, CPLD and RTL8211F BMC Ethernet PHY chip.

7.10.2. Reset

Several reset signals are connected between Ampere® processor, BMC, and PCIE devices on Mt. Mitchell MB:

- System Reset button: This is a cold reboot of the mainboard with no effect on BMC's functionality. Upon receiving the System Reset signal from the user, BMC resets the mainboard.
- BMC Reset button: Resets the BMC only, with no effect on the mainboard.
- Ampere® does not have a dedicated PERST for each controller. The solution is to use CPU I2C7 which is connected to IO expanders
- Ampere® CPU supports cold reset and warm reset. They are handled by Mpro
- After all power rails are stable, SYS_RESETN must be held LOW for a minimum of 10 µs.

Figure 22 shows the reset signal connections among Ampere® processor, BMC, and CPLD.

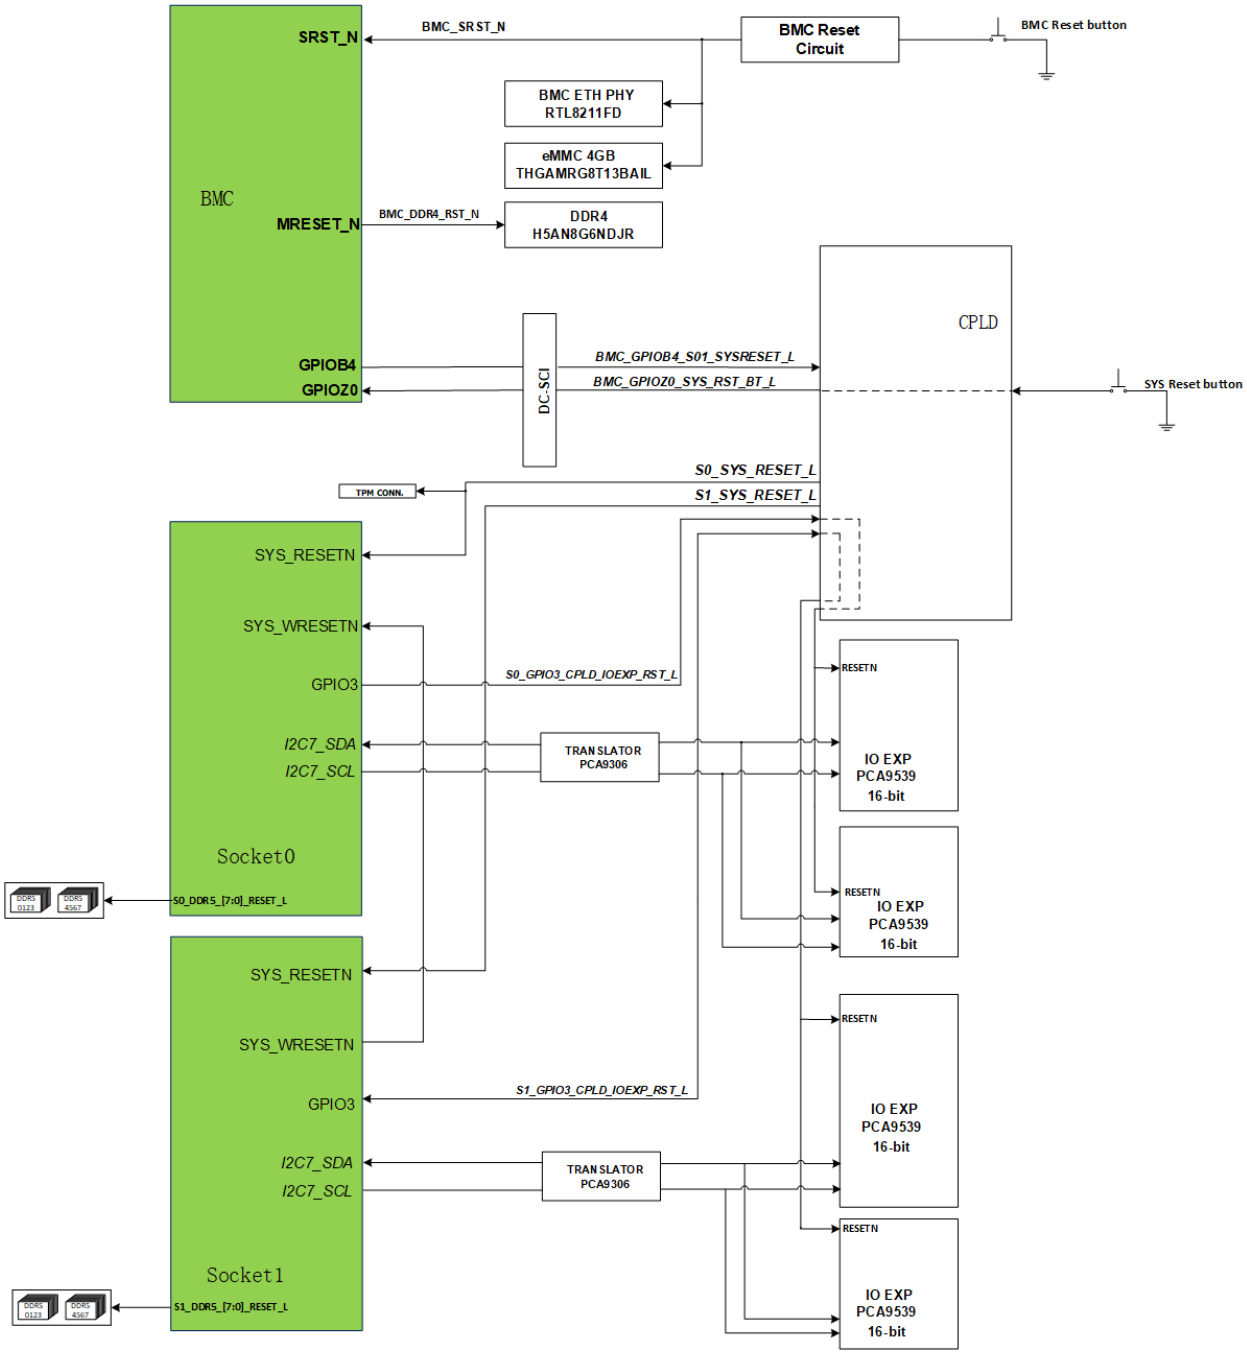

Figure 22 Mt. Mitchell Reset Signals

7.10.3. Power-on Sequence

Figure 23 shows the Ampere® processor power-on sequence for 2 socket platform where I2C9_ALERT9_N signal, which is supposed to trigger the execution of socket1 FW, is directly connected between processors in master/slave sockets.

First, socket0 (master socket) power rails will be turned on in order of P0V85_SOC, P0V85_SOC_RC_DDR0/1, P0V9_RC5A_AVDD, P1V8_RC5A_AVDDH, P1V1_VDDQ_0123/4567, P1V8_SOC, P1V2_SOC. The SOC_PWRGD of socket0, S0_SOC_PWRGD, will be asserted. The same power-on sequence should be applied to socket1 (slave socket).

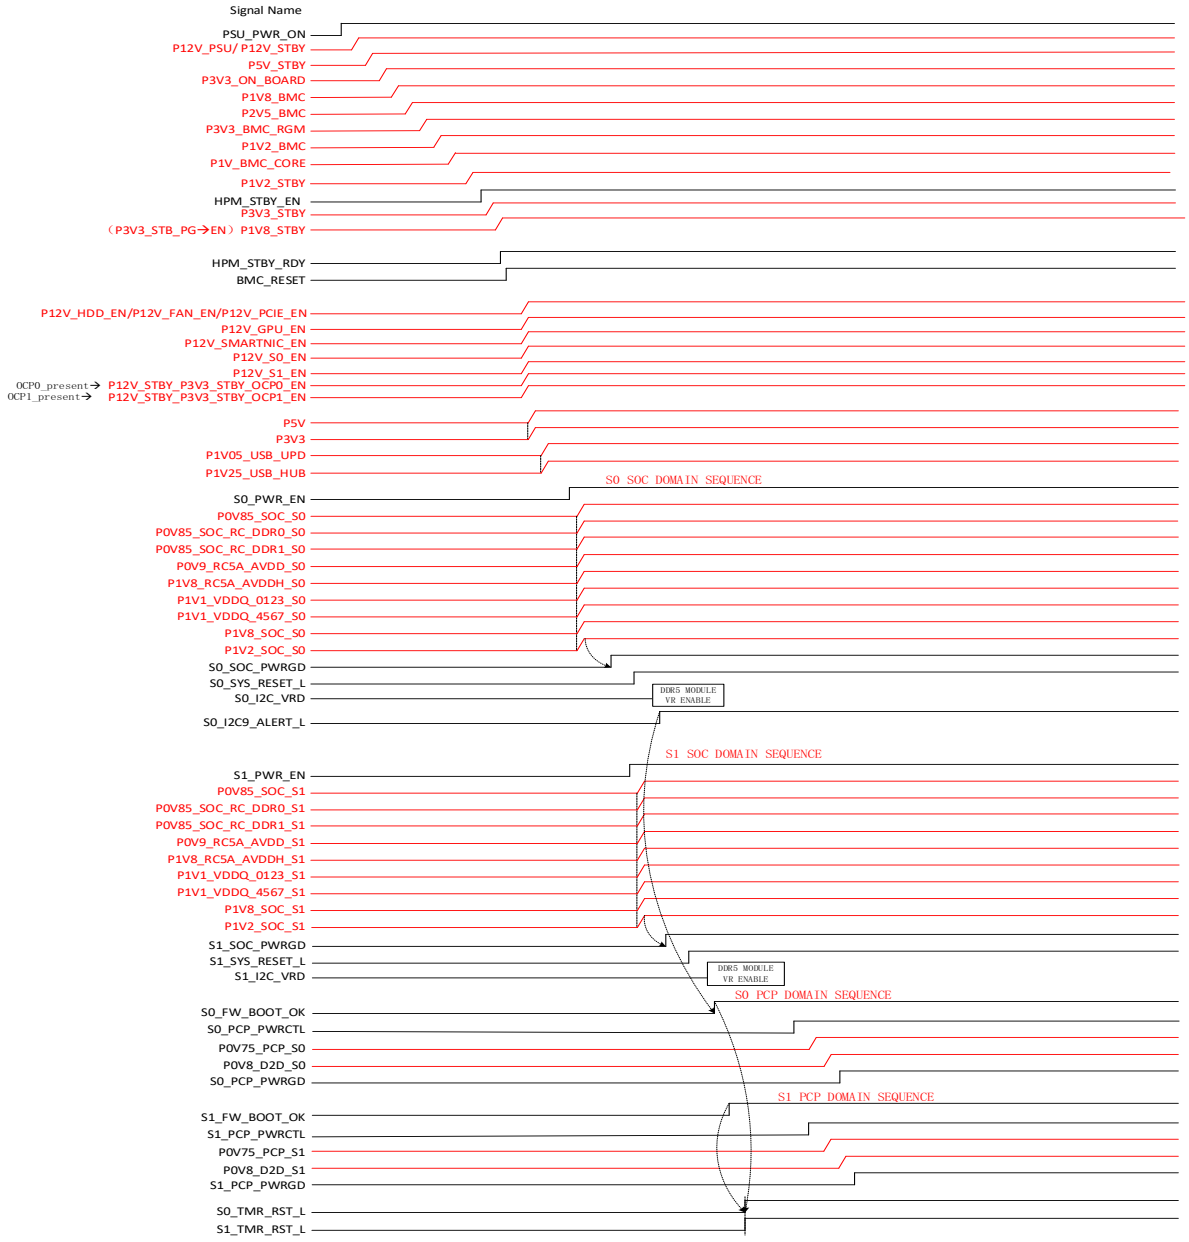

Figure 23 Power-on Sequence

7.11. UART

Mt. Mitchell contains five UARTs (1V8 level), whose assigned functions are listed below:

- UART0 is BIOS and OS boot console.
- UART1 is Mpro console is required for debugging Mpro.

- UART2 is ATF console.
- UART4 is SECpro console.
- UART3 is not used.

Below figure shows the connections of UART.

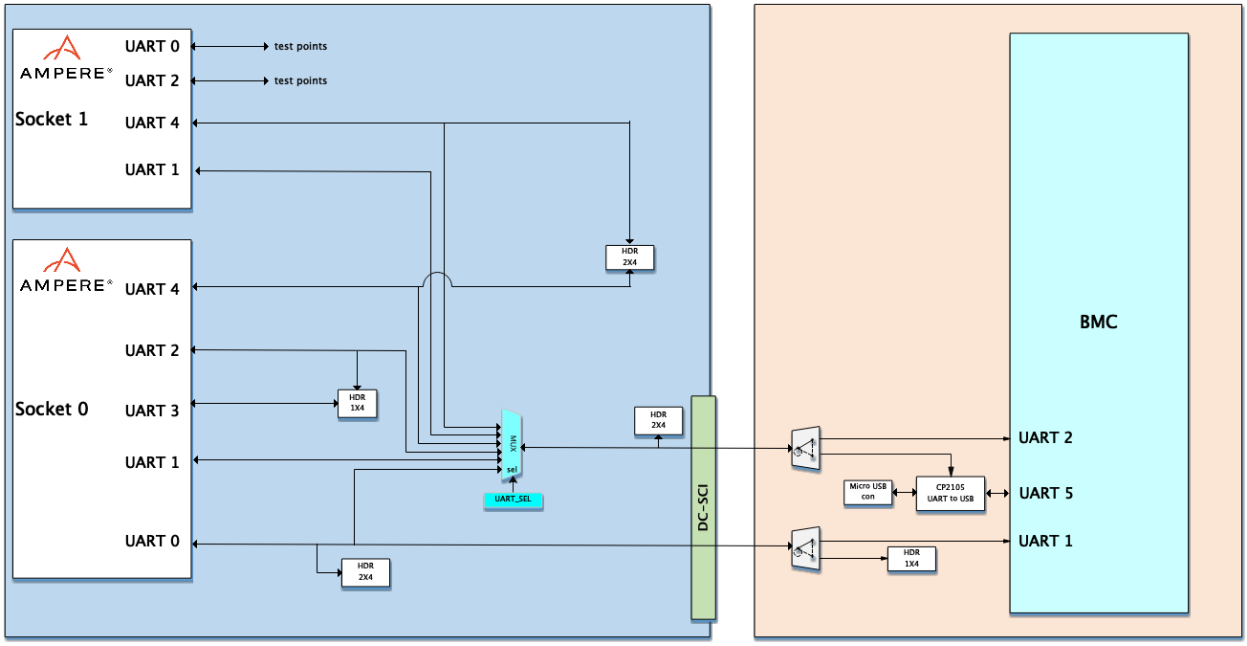

Figure 24 UART for Ampere® CPU and BMC

8. Thermal Design Requirements

Table 14 lists the thermal resistance values for the package along with the maximum operating junction temperature, the thermal throttling and shutdown temperature for Ampere@ processor. The θ_{JC}, T_{CJ}, TM1, and TM2 values in the table are varied depending on SKU.

Parameter	Symbol	Ampere@ CPU
Maximum continuous operating junction temperature[1]	T_{CJ}	100 °C
SoC thermal throttling temperature[1]	TM1	105 °C
SoC thermal shut down temperature[1]	TM2	120 °C
Storage temperature range[2]	T_{STG}	-55 °C to +150 °C
Operating junction temperature range[2]	T_J	0 °C to +125 °C

Notes:
1. TCJ as specified in this table is for use as a system thermal design guideline. The SoC can run continuously at this temperature. To achieve maximum performance, ensure that the thermal solution can maintain TCJ at this value.
In case of operational or cooling failures, the following built-in mechanism in firmware will protect the SoC from a thermal runaway or getting damaged:
– HIGHTEMP_N (TM1): The SoC temperature at which thermal throttling will be triggered. The CPU frequency is throttled down in steps of 50 MHz.
– OVERTEMP_N (TM2): The SoC temperature at which a shutdown will be triggered. The entire SoC is powered off under this condition.
2. This value is not a specification of the operational temperature range; it is a stress rating only.

Table 14 Ampere® CPU Thermal Specification

9. I/O System

9.1. Front Panel Drive Bays and Controls

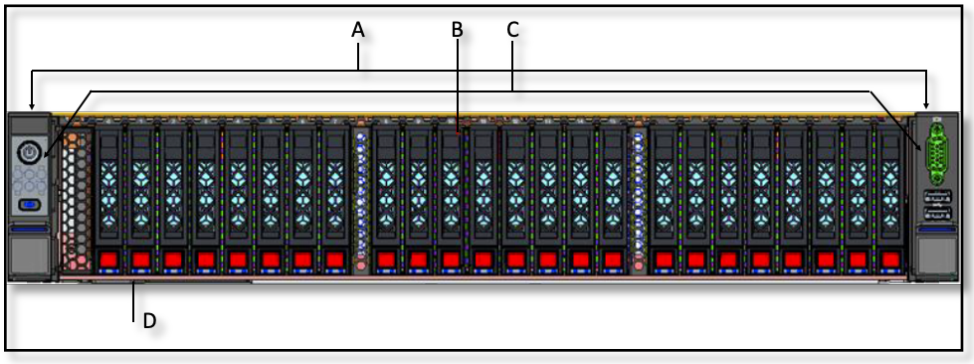

Figure 25 Mt. Mitchell 2U Front Drive Bays

Item	Description
A	Front Rack Handles
B	Storage Drive Bay
C	Front Control Panel
D	Information Tag

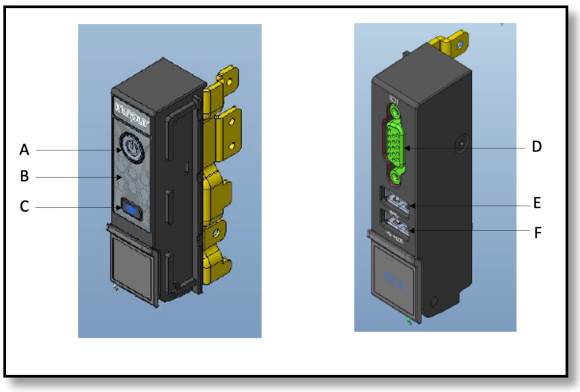

Figure 26 Mt. Mitchell 2U Front Panel

Item	Description
A	Power button with LED
B	System Status/Fault LEDs
C	ID button with LED
D	VGA
E	USB 3.0
F	USB 2.0

9.2. OCP V3.0 NIC and Single-Host/Multi-host Configurations

Supports up to two OCP NIC 3.0 (SFF) x16 cards in the rear, one from each SoC. Dual OCP NIC Single-host and Multi-host configurations are shown in below figure.

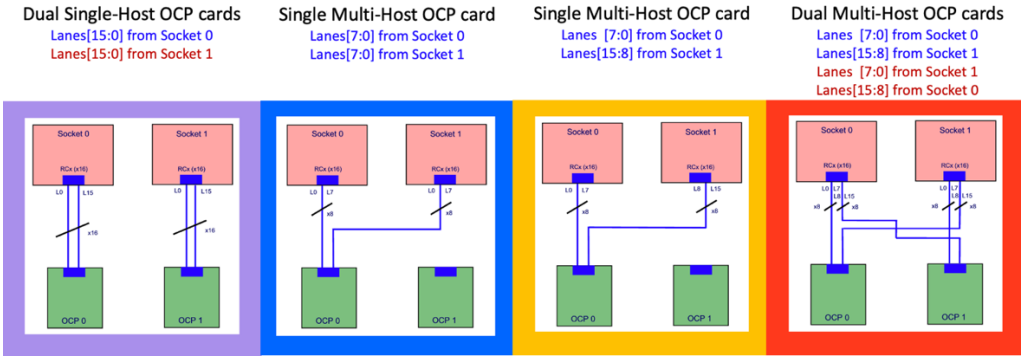

Figure 27 Mt. Mitchell OCP Single-host/Multi-host Configurations

9.3. M.2 Configuration

M.2 riser plugged into Mainboard can support up to 2x M.2 slots:

- M.2 slots on M.2 riser support 80/110mm M.2 NVMe x4 PCIe lanes
- These slots are connected CPU0's PCIe RC1[0:7] for OS booting.

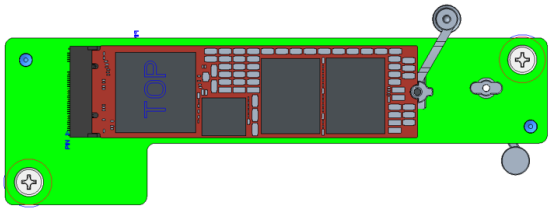

Figure 28 Mt. Mitchell M.2 Configuration

10. Rear Side Power, I/O, Expansion Board and Midplane Subsystems

Figure 29 shows the rear side power and I/O when Mt. Mitchell MB is installed in 2U chassis and Table 15 lists the components in each location.
Figure 30 shows the system management module and Table 16 lists the components in each location.

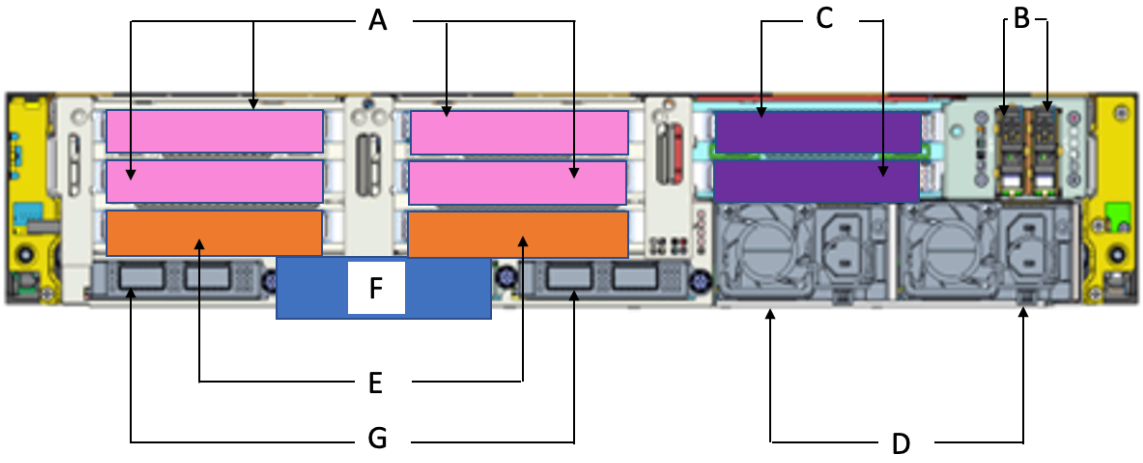

Figure 29 Mt. Mitchell 2U Rear Side I/O

Item	Description
A	FHFL x16 PCIe Cards
B	E1.S
C	FHFL x16 PCIe Cards
D	Power Supply Unit (PSU)
E	FHHL x16 PCIe Cards
F	System Management Module
G	OCP Expansion Slots

Table 17 Mt. Mitchell 2U Rear Side I/O

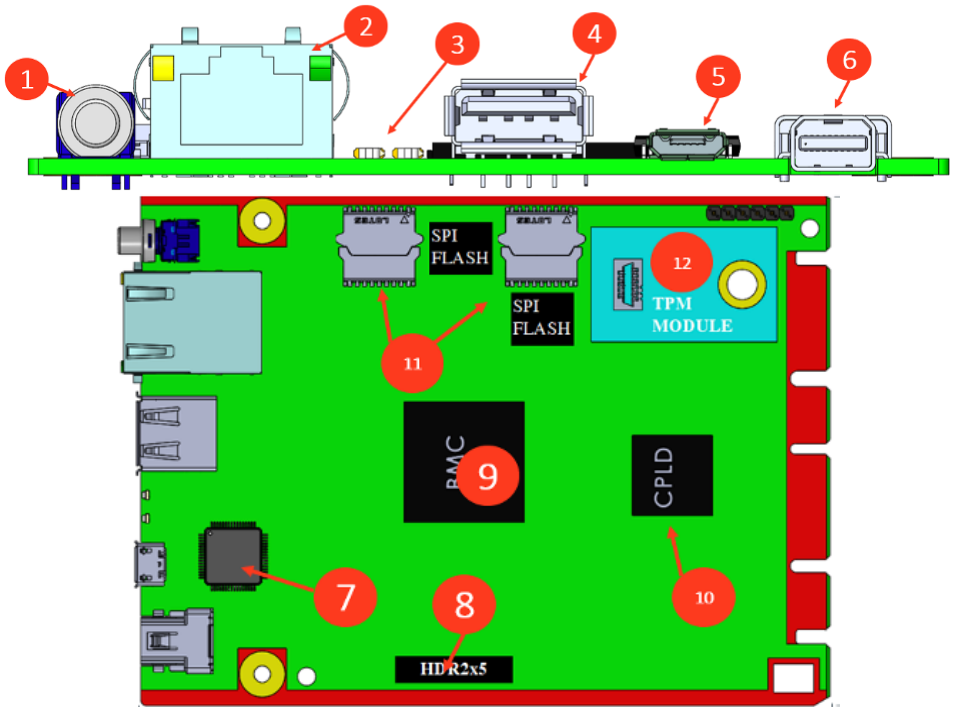

Figure 30 Mt. Mitchell System Management Module (DC-SCM)

Item	Description
1	UID Button with LED
2	RJ45 Management Port Micro USB 2.0 (for BMC UART)
3	Power LED Status LED
4	USB 2.0
5	Micro USB 2.0 (for BMC UART)
6	Mini DP
7	Cypress UART-USB Bridge Chip
8	Header 2x5 for external VGA port
9	BMC AST2600
10	CPLD
11	Flash & Sockets
12	TPM Module

Table 18 Mt. Mitchell System Management Module

11. Mechanical

Highlights of Mt. Mitchell in mechanical point of view.

- CPU and heat sink
  - CPU socket is soldered on Mt. Mitchell motherboard to hold CPU
  - 2U heat sink can be mounted on top of the socket
- DIMMs
  - Standard DDR5 DIMM connectors
- OCPv3 NIC
  - Two OCPv3 NIC cards along with 4C+ connector
  - Chassis supports latch or screw based
- PCIE riser cards
  - There are cages with latch and screw to hold the riser cards
- FAN
  - Two sets of FAN connectors which are aligned with hot-swappable FAN cage connectors
- CPRS PSU
  - Two connectors are located at left side
  - PSU is supported by latch on chassis
- System Management Module
  - DC-SCM compliant system management module.

12. Onboard Power System

12.1. Mt. Mitchell Power Delivery Block Diagram

Mt. Mitchell power system is designed to support the Ampere® processor, BMC sub-system including fans, and the PCIe devices which are connected to PCIe riser cards, OCPv3 connectors, and M.2 connectors. The Figure 31 shows the detail connections of Mt. Mitchell power rails.

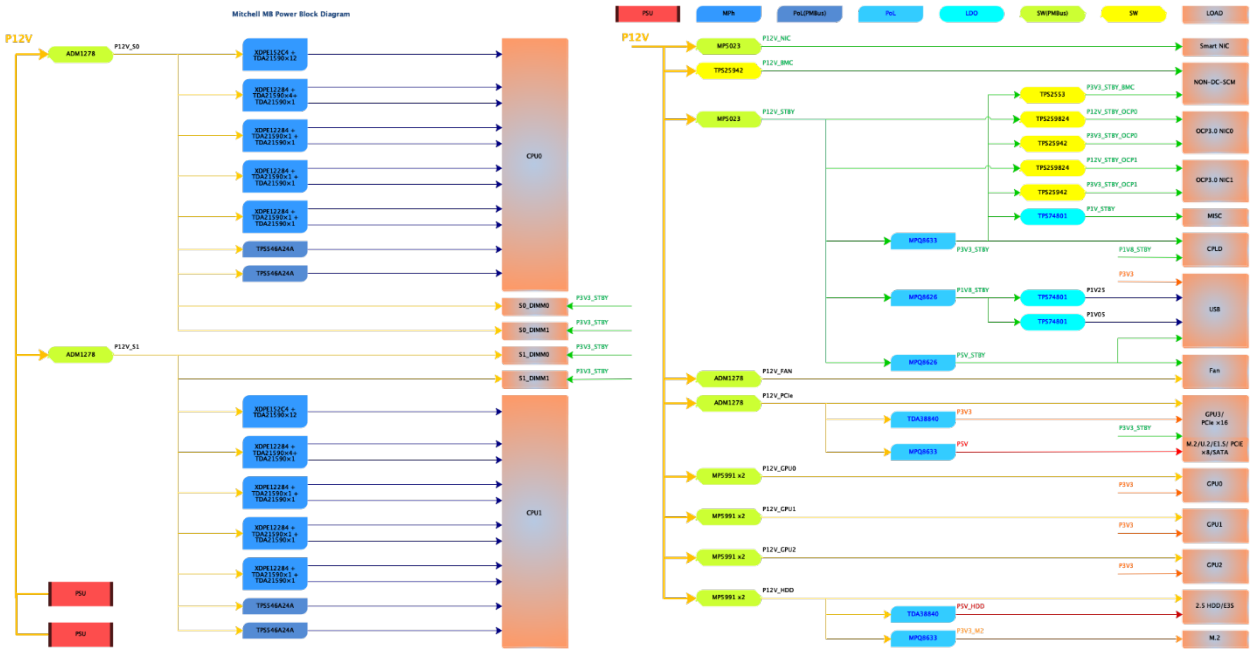

Figure 31 Mt. Mitchell Power Delivery Block Diagram

There are two high power domains for Ampere® processor:

- PCP power domain for CPU cores and mesh interconnects
- SoC power domain for SECpro, Mpro and AHBC

Infineon voltage regulators, XDPE152C4D and XDPE12284C, are used for these high-power domains along with TDA21590 power stages. Both VRD's are PMBus™ system interface rev 1.2 compliant for telemetry of voltage, current, power, temperature, and fault conditions. Also, they support dynamic output voltage transitions with programmable slew rates via PMBus interface with 100kHz ~ 400kHz frequency.

For other power rails of the Ampere® processor, TI TPS546A24 voltage converter is used.

PCIE peripherals need 3.3V and 12V power. The 3.3V power is converted from 12V main power by Infineon TDA38840 to support up to 32A. The PCIE 12V power is directly connected from power supply through eFuse ADM1278.

13. Environmental Regulations/Environmental Requirements

The motherboard shall meet the following environmental requirements:

- NA

The full system shall meet the following environmental requirements:

- NA

Mt. Mitchell shall meet the technical requirements in the following EMC, safety and environmental compliance standards.

- NA

14. Prescribed Materials

NA.

15. Software Support

NA.

16. System Firmware

All products seeking OCP Accepted™ Product Recognition must complete the Open System Firmware (OSF) Tab in the 2021 Supplier Requirements Checklist.

If based on an open bios (like AMI's Aptio-OpenEdition), a completed checklist shall be uploaded and made available on the OCP Github.
https://github.com/opencomputeproject/Hardware-Management/amperecomputing/

17. Hardware Management

17.1 Compliance

All products seeking OCP Inspired™ or OCP Accepted ™ Product Recognition shall comply with the OCP Hardware Management Baseline Profile V1.0 and provide such evidence by completing the Hardware Management Tab in the 2021 Supplier Requirements Checklist.

17.2 BMC Source Availability (if applicable)

All Products seeking OCP Accepted™ Product Recognition shall have source code and binary blobs submitted for BMC, if applicable.

The BMC management source code shall be uploaded at: https://github.com/opencomputeproject/Hardware-Management/amperecomputing/mtmtchell

If the BMC is based on an open source BMC (like AMI's MegaRAC-OpenEdition), the BMC source code shall be uploaded and made available on the OCP Github.

18. Security

All products seeking OCP Inspired™ or OCP Accepted™ Product Recognition shall have a completed Security Profile in the 2021 Supplier Requirements Checklist. Whether the answer is a yes or no, the profile must be completed. For Additional Security Badges (Bronze/Silver/Gold), please fill out the Security Profile in accordance with the requirements for that level. Security Badges will be reassessed on an annual basis as requirements are subject to change.

19. References

[1] “Title”, publication year, publication journal/conference/standard, volume, pages, link to publication if available
[2] OCP Profiles - https://github.com/opencomputeproject/OCP-Profiles
[3] Redfish Interop Validator - https://github.com/DMTF/Redfish-Interop-Validator
[4] Redfish Service Validator - https://github.com/DMTF/Redfish-Service-Validator
[5] Redfish Service Conformance Check - https://github.com/DMTF/Redfish-Service-Conformance-Check

Appendix A - Checklist for IC Approval of this Specification (to be completed by Contributor(s) of this Spec)

Complete all the checklist items in the table with links to the section where it is described in this spec or an external document.

Item	Status or Details	Link to detailed explanation
Is this contribution entered into the OCP Contribution Portal?	Yes	
Was it approved in the OCP Contribution Portal?	Yes	
Is there a Supplier(s) that is building a product based on this Spec? (Supplier must be an OCP Solution Provider)	Yes	Inspur Corporation
Will Supplier(s) have the product available for GENERAL AVAILABILITY within 120 days?	Yes	Please have each Supplier fill out Appendix B.

Appendix B - Inspur - OCP Supplier Information and Hardware Product Recognition Checklist

Company: Inspur Corporation
Contact Info:

Product Name: Mt. Mitchell
Product SKU#: TBD
Link to Product Landing Page: TBD

The following is needed for OCP hardware product recognition:

For OCP Inspired™

- All Suppliers must be a Silver, Gold or Platinum Member.
- Declare product is 100% compliant with specification.
- Complete the OCP Inspired™ Product Recognition Checklist, which includes hardware management conformance checks and security profile.

For OCP Accepted™

- All Suppliers must be an OCP Member. All corporate membership levels are eligible.
- Complete the OCP Accepted™ Product Recognition Checklist, which includes hardware management conformance checks, security profile and open system firmware conformance checks.
- Submit a design package meeting OCP Hardware Design Guideline Contribution Checklist (if not already submitted by the contributor). If already submitted, declare the product is 100% compliant with the design package.
- Submit a firmware package including a firmware image, build scripts, documentation, test results and a tool that verifies modifications
- Submit the BMC source code, if applicable to product type

Please complete the OCP Inspired™ Product Recognition Submission Checklist or OCP Accepted™ Product Recognition Checklist and the following table.

Requirements	Details	Links
Which Product recognition?	OCP Accepted™	Provide link for the appropriate Product Checklist
If OCP Accepted™, who provided the Design Package?	Inspur Corporation	Link to OCP Contribution Database

Where can a potential adopter purchase the product?		Link to OCP Marketplace

Appendix C - Contribution Process FAQs

As a contributor to a hardware specification, here are some questions that often come up.

Q1. What type of hardware specification am I contributing to OCP? Is it any of the below?

a. base specification for a de-facto standard (new standard with no hardware product on the horizon)
b. base specification for an intended physical <hardware product type> (product may be coming but within the next 1-2 years)
c. modification of an existing specification (state which existing spec is being modified)
    i. either a complete revision update or
    ii. a minor version update
d. design spec (based on an existing base specification) with more refined design details (product coming in 12-15months)
e. a detailed specification for a <hardware product type> for a very specific product being available in 3-6months of approval of this Spec
f. If none of the above, please contact OCP Staff for better direction.

Q2. How do I know if what I am contributing will be accepted by OCP?

a. Before contributing any specifications, please contact either OCP Staff (Archna Haylock or Michael Schill) or the Project Lead for the Project that best represents your contribution. For example, if you are contributing a Server Specification, please contact one of the Server Project Leads. You can see all the Projects here.
b. They will help you with your contribution and help you navigate the process.

Q3. What is the contribution process for my hardware spec?

a Follow the flow for your spec type here.
b This flow is subject to change so please check with the OCP Staff for more information or any questions.

Q4. What if my spec is not developed yet and I want to collaborate with other companies?

a. Please contact either OCP Staff (Archna Haylock or Michael Schill) or the Project Lead for the Project that best represents your contribution.
b. They will help you find other collaborators and help you with the contribution process for a multi-party contribution.

Q5. I have a question on the Contribution License Agreement.

a Please contact OCP Staff and we can help you with questions.

Q6. Do I need to have a product in order to contribute a spec?

a. Please Please see Q1. Some types of specs do not require an immediate product. Some do. Please work with the OCP Staff on better direction on your specification type.