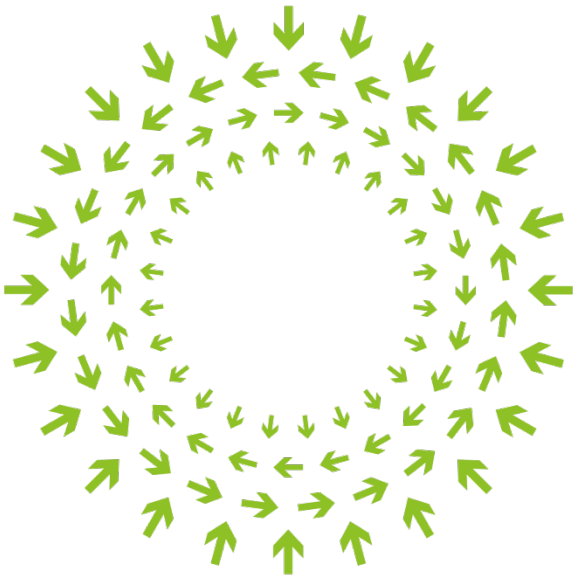

OPEN

Compute Project

1S Server Design Specification

V0.4

Author: Yan Zhao, Hardware Engineer, Facebook

1 Scope

This specification describes the design of the 1S server.

2 Contents

3 Overview

This document describes a single socket server design, which is referred to hereinafter as a 1S server. The 1S server uses a server class System-on-a-Chip (SoC), which usually integrates a multi-core CPU, a memory controller, and input/output (IO) devices in one package. With the higher performance and the higher IO density of the multi-core SoCs, a 1S server is perfect for single or multi-node platforms where performance and IO density are the key factors. However, platform designers must provide adequate power and cooling to properly handle the SoC's power and thermal requirements.

Figure 3-1 illustrates a general 1S server block diagram.

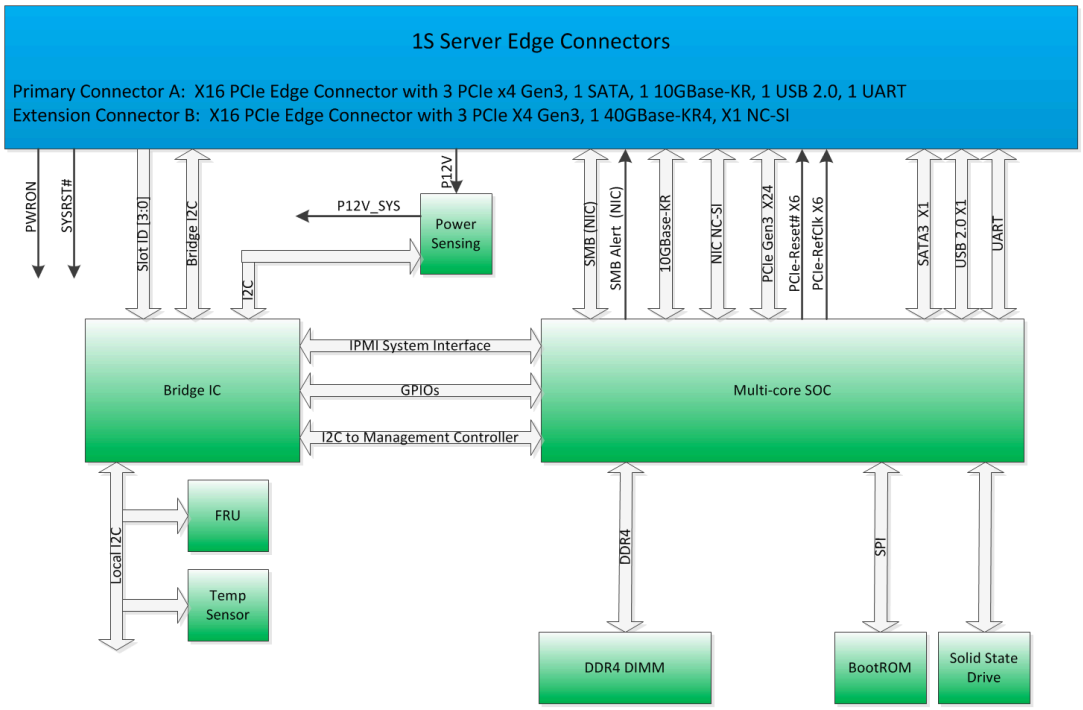

Figure 3-1 1S Server Block Diagram

Compared to the Open Compute Project (OCP) micro-server (hardware specification available at http://files.opencompute.org/oc/public.php?service=files&t=0ab1c41dd11111c4f3ceafc175883fcb), the 1S server is significantly larger to accommodate the larger SoC, more DIMM slots, and its higher power consumption. The 1S server specification defines two x16 PCIe edge connectors as the interface to the platform that hosts the 1S server. The primary x16 PCIe edge connector is mandatory. It supports:

- x12 PCIe lanes in a maximum of three ports
- A 10GBase-KR
- A SATA port
- A USB port
- A Universal Asynchronous Receiver/Transmitter (UART)
- An I^2C bus for server management

- A System Management Bus (SMBus) as the sideband of the integrated network controller

An optional extension of an x16 PCIe edge connector can be used to provide x12 extra PCIe lanes in a maximum of three ports, a 40GBase-KR4 port, and a Network Controller Sideband Interface (NC-SI) as an alternative high-speed sideband option of the integrated network controller.

The 1S server shall have an on-card, high-performance storage component such as a solid state drive (SSD) or similar for the boot device and log storage. It is required to support at least one SATA and/or PCIe based SSD drive in the M.2 form factor. The minimum capacity recommended is 128GB. Because the 1S server supports SATA and PCIe ports on the edge connectors, it is possible to add more storage on the platform.

The 1S server can use either an external network interface controller on the platform through its PCIe interface or the SoC's integrated network controller. The network controller must be either the 10Gbase-KR or 40Gbase-KR4 type Ethernet controllers. Generally, the 1S server assumes a Baseboard Management Controller (BMC) on the platform side. When the integrated network controller on the SoC is used as a shared NIC, it shall support a NC-SI or SMBus as a sideband to the BMC on the platform in all power states.

The 1S server receives 12.5V from the platform with a maximum current load of 7.7A from the primary edge connector. It receives an additional 7.7A from the extension edge connector. Thus, in theory, a 1S server's maximum power consumption is 192W. However, the platform defines and controls the maximum power used by the 1S server. The 1S server shall use a current sensor at the power input in order to accurately measure the power consumption of the entire server for power management purposes. The 1S server shall support an Advanced Configuration Power Interface (ACPI)-compliant power button and reset signals from the platform.

A Bridge IC is used to provide a common management interface for the 1S server to enable SoC-agnostic designs. The Bridge IC is the management controller on the 1S server and the bridge between the BMC and SoC. The Bridge IC manages the 1S server on behalf of the BMC on the platform and bridges the BMC and SoC's own internal management controller (if there is one). To maximize the communication bandwidth between the BMC and the Bridge IC, a dedicated point-to-point I^2C bus is used as the interface.

The 1S server's Field Replaceable Unit (FRU) EEPOM and thermal sensors are connected to the Bridge IC 's other I^2C buses. There are multiple GPIOs between the Bridge IC and the SoC for error reporting and other management purposes. If the SoC has an internal management microcontroller, it should be accessible to the Bridge IC via an I^2C interface. An IPMI system interface between the Bridge IC and SoC, such as a KCS (Keyboard Controller Style) or SSIF (SMBus System Interface), shall be provided to enable in-band communications. The BMC shall access the 1S server's thermal sensor, FRU, SoC GPIOs, and internal management controller via the Bridge IC with standard IPMI commands.

The BMC and Bridge IC shall provide mechanisms to update the 1S server's programmable devices, such as the SoC's boot ROM, the Bridge IC's boot ROM, and so on. This enables the user to diagnose and fix 1S servers remotely; greatly improving serviceability and availability.

4 License

As of December 19, 2014, the following persons or entities have made this Specification available under the Open Web Foundation Final Specification Agreement (OWFa 1.0), which is available at http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owfa-1-0:

Facebook, Inc.

You can review the signed copies of the Open Web Foundation Agreement Version 1.0 for this Specification at http://opencompute.org/licensing/, which may also include additional parties to those listed above.

Your use of this Specification may be subject to other third party rights. THIS SPECIFICATION IS PROVIDED "AS IS." The contributors 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 Specification. The entire risk as to implementing or otherwise using the Specification is assumed by the Specification implementer and user. 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 SPECIFICATION OR ITS GOVERNING AGREEMENT, 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.

5 Mechanical

5.1 Mechanical Outline

The 1S server card must have an x16 PCIe edge connector. An additional extension x16 PCIe edge connector is optional.

The overall dimentions of the general card are 210mmx110mm or 210mmx160mm. See Figure 5-1 for the specificaiton drawing inlcuding keep-out zones and component restrictions.

The thickness of the PCB shall be 1.57±0.16mm to accommodate regular PCIe x16 connectors on platforms. Components of significant height will be placed on the A-side with maximum height limit of 34.20mm. Low profile components of height less than 2.67mm can be placed on the B-side. The heat sink fins will be parallel to the long edge of the card.

The component and trace free areas for the guide features are 5mm keep-out zones on the sides of the card. This area will be silk-screened white on both sides.

5.2 PCIe Edge Connector

The key dimensions, edge chamfer, pad layout (including a shorter pad for PRSNT# signal), placement, and dimensions of the card edge connector match the PCI Express Card electromechanical specification.

The GND planes underneath the pads for the edge connector on the card must be recessed according to the PCI Express Card electromechanical specification to improve signal integrity.

5.3 Platform Design

Platform design details are not discussed in this specification.

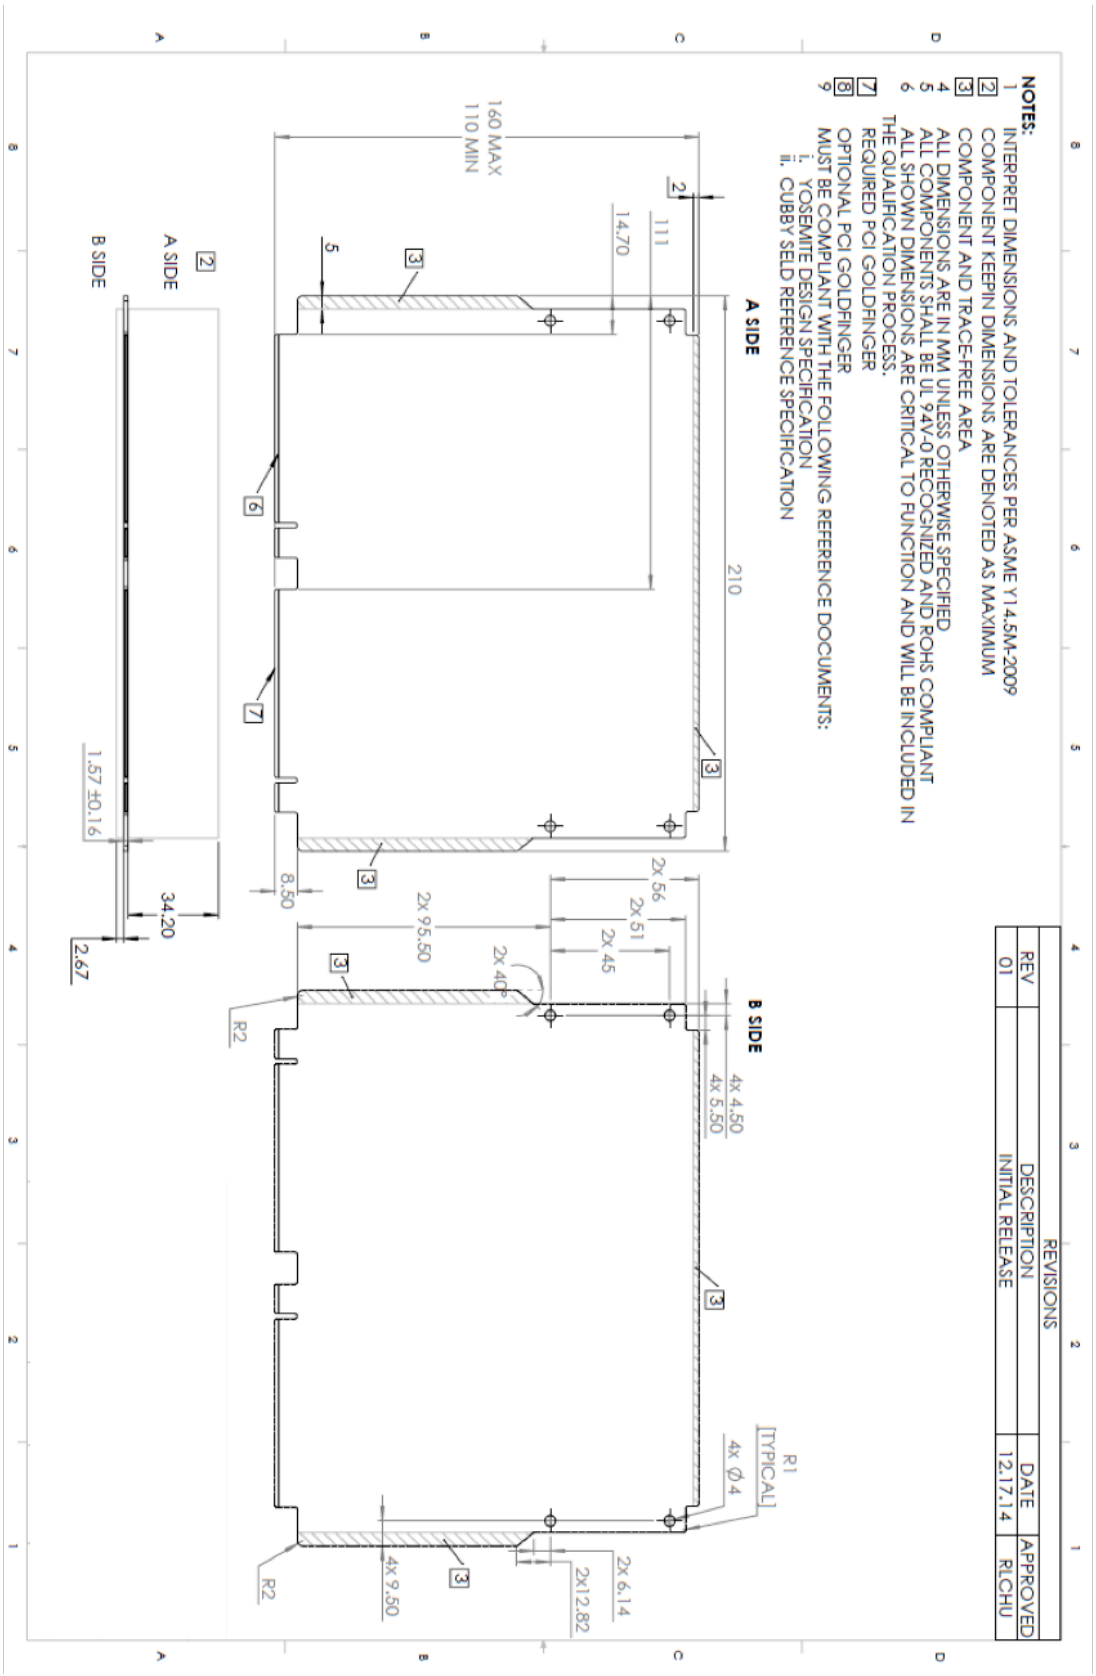

Figure 5-1 1S Server Mechanical Drawing

6 Thermal

6.1 Data Center Environmental Conditions

6.1.1 Location of Data Center/Altitude

The data center may be located at a maximum of 1000m above sea level. Any variation of air properties or environmental differences due to the high altitude must be included in the thermal design.

6.1.2 Cold-Aisle Pressurization

The data center will maintain the cold aisle pressure between 0"H_2O and 0.01"H_2O. The thermal solution of the system accommodates the worst-case operational pressurization in the data centers, which is 0"H_2O.

6.1.3 Relative Humidity

The data center will maintain the relative humidity between 10% and 90%. In the thermal design, the environmental condition changes due to the high altitude may not be considered when the thermal design can meet the requirement with the maximum relative humidity (90%).

6.2 Server Operational Conditions

6.2.1 System Volumetric Flow or Linear Feet per Minute

The unit of airflow (or volumetric flow) used for this specification is cubic feet per minute (CFM). The maximum allowable airflow per watt in the system must be 0.15. The desired airflow per watt is 0.12 or lower up to 45°C (113°F) ambient temperature. The linear feet per minute (LFM) should be less than 470.

6.2.2 Thermal Margin

The thermal margin is the difference between the maximum theoretical safe temperature and the actual temperature. The board design operates at an inlet temperature of 45°C (113°F) outside of the system with a minimum 5% thermal margin for every component on the card.

6.2.3 Upper Critical Threshold

The upper critical threshold (UCT) setting should allow the detecting of abnormal thermal behaviors in the system. The UCT values for the sensors that are not used in Fan Speed Control (FSC) should use 15% thermal margin from the worst experiment data. The UCT values for the sensors used in FSC, except for CPU, inlet, and outlet sensors, should use 20% thermal margin from the worst experiment data.

6.2.4 Thermal Testing

Thermal testing must be performed at a low inlet temperature 15°C (59°F) and up to 50°C (122°F) or higher inlet temperature to guarantee the design is free of thermal defect and has high temperature reliability.

6.3 Heat Sink Requirements

The heat sink must be a thermally optimized design at the lowest cost. Heat sink installation must be uncomplicated. Passive cooling is desired. Heat sinks must not block debug headers or connectors. The heat sink fins should be aligned with the airflow direction, which is shown in Figure 3-1. The LFM for the heat sink should be less than 370, or the pressure drop across the heat sink should be less than 0.045 inches H_2O.

6.4 Temperature and Power Sensors

Each card must provide:

- Temperature sensors for the Soc
- Power reading for the SoC
- SO-DIMMs
- Memories (if they are used)
- Voltage regulators
- One inlet ambient temperature sensor
- Any other critical chipsets

Each sensor's temperature and power readings must be readable via the management sideband interface to the platform. Additionally, over-temperature thresholds must be configurable and an alert mechanism must be provided to enable thermal shutdown and/or an increase in airflow. The sensors are accurate to +/-2°C and desired to be within 2% tolerance across whole operation temperature range.

Two ambient temperature sensors are placed along the edges of the card on the A-side. Figure 6-1 the desired areas for ambient temperature sensors are highlighted.

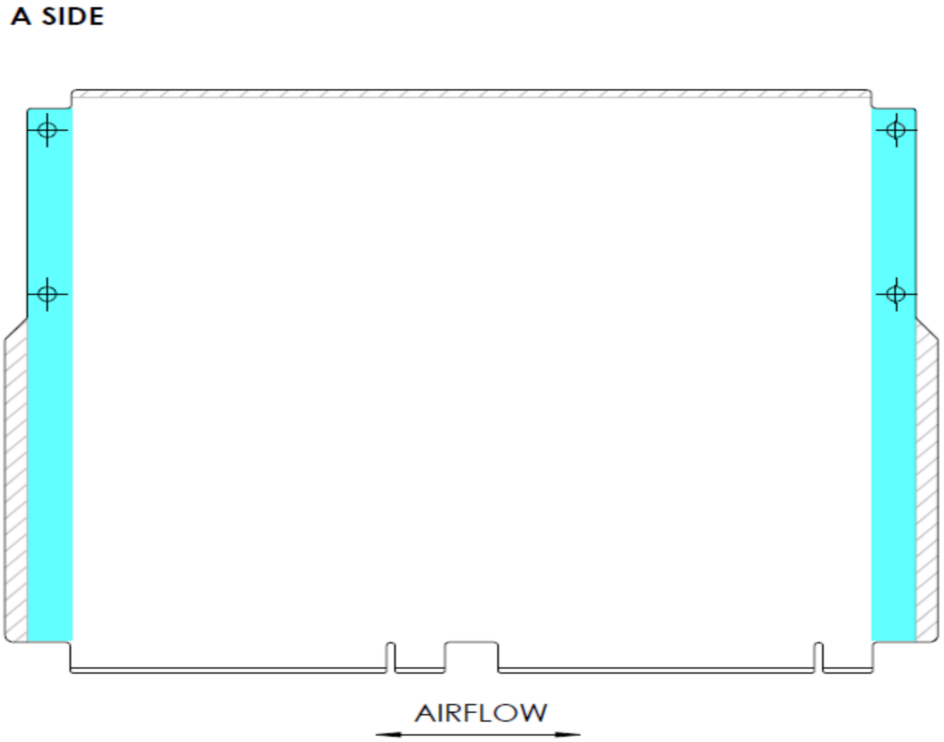

Figure 6-1 Ambient Temperature Sensor Location (Circled Area in Blue)

7 Electrical

7.1 Design Guidelines

Refer to the SoC vendor's documents for design guidelines.

7.2 Primary X16 Edge Connector A

The 1S server uses the x16 edge connector pin assignments as defined in the OCP micro-server specification with a few exceptions, such as newly defined FAST_THROTTLE_N and 10GBase-KR signals. This primary X16 Edge connector is referred as Connector A. This connector is mandatory.

The OCP Micro Server hardware spec V0.7 is available at: http://files.opencompute.org/oc/public.php?service=files&t=0ab1c41dd11111c4f3ceafc175883fcb

Table 1: 1S Server Primary X16 OCP Edge Connector A Pin-Out

Default Pin-Out			
Pin Name	**B Side**	**A Side**	**Pin Name**
P12V	1	1	PRSNT_A#
P12V	2	2	P12V
P12V	3	3	P12V
GND	4	4	GND
I2C_SCL	5	5	SVR_ID0/GPIO0
I2C_DATA	6	6	SVR_ID1/GPIO1
GND	7	7	COM_TX
PWR_BTN#	8	8	COM_RX
USB_P	9	9	SVR_ID2/GPIO2
USB_N	10	10	SVR_ID3/GPIO3
SYS_RESET#	11	11	PCIE0_RESET#
I2C_ALERT#	12	12	GND
GND	13	13	PCIE0_REFCLK_P
GND	14	14	PCIE0_REFCLK_N
PCIE0_TX0_P	15	15	GND
PCIE0_TX0_N	16	16	GND
GND	17	17	PCIE0_RX0_P
GND	18	18	PCIE0_RX0_N

PCIE0_TX1_P	19	19	GND
PCIE0_TX1_N	20	20	GND
GND	21	21	PCIE0_RX1_P
GND	22	22	PCIE0_RX1_N
PCIE0_TX2_P	23	23	GND
PCIE0_TX2_N	24	24	GND
GND	25	25	PCIE0_RX2_P
GND	26	26	PCIE0_RX2_N
PCIE0_TX3_P	27	27	GND
PCIE0_TX3_N	28	28	GND
GND	29	29	PCIE0_RX3_P
GND	30	30	PCIE0_RX3_N
SATA0_TX_P	31	31	GND
SATA0_TX_N	32	32	GND
GND	33	33	SATA0_RX_P
GND	34	34	SATA0_RX_N
PCIE1_REFCLK_P	35	35	GND
PCIE1_REFCLK_N	36	36	GND
GND	37	37	PCIE2_REFCLK_P
GND	38	38	PCIE2_REFCLK_N
PCIE1_RESET#	39	39	GND
PCIE2_RESET#	40	40	GND
GND	41	41	FAST_THROTTLE_N
GND	42	42	NIC_SMBUS_ALERT#
NIC_SMBUS_SCL	43	43	GND
NIC_SMBUS_SDA	44	44	GND
GND	45	45	KR_RX_P
GND	46	46	KR_RX_N
KR_TX_P	47	47	GND
KR_TX_N	48	48	GND
GND	49	49	PCIE1_RX0_P
GND	50	50	PCIE1_RX0_N
PCIE1_TX0_P	51	51	GND

PCIE1_TX0_N	52	52	GND
GND	53	53	PCIE1_RX1_P
GND	54	54	PCIE1_RX1_N
PCIE1_TX1_P	55	55	GND
PCIE1_TX1_N	56	56	GND
GND	57	57	PCIE1_RX2_P
GND	58	58	PCIE1_RX2_N
PCIE1_TX2_P	59	59	GND
PCIE1_TX2_N	60	60	GND
GND	61	61	PCIE1_RX3_P
GND	62	62	PCIE1_RX3_N
PCIE1_TX3_P	63	63	GND
PCIE1_TX3_N	64	64	GND
GND	65	65	PCIE2_RX0_P
GND	66	66	PCIE2_RX0_N
PCIE2_TX0_P	67	67	GND
PCIE2_TX0_N	68	68	GND
GND	69	69	PCIE2_RX1_P
GND	70	70	PCIE2_RX1_N
PCIE2_TX1_P	71	71	GND
PCIE2_TX1_N	72	72	GND
GND	73	73	PCIE2_RX2_P
GND	74	74	PCIE2_RX2_N
PCIE2_TX2_P	75	75	GND
PCIE2_TX2_N	76	76	GND
GND	77	77	PCIE2_RX3_P
GND	78	78	PCIE2_RX3_N
PCIE2_TX3_P	79	79	GND
PCIE2_TX3_N	80	80	GND
GND	81	81	P12V
GND	82	82	P12V

7.3 Extension X16 Edge Connector B

The 1S server also implements an extension x16 edge connector to bring out additional x12 PCIe lanes, one 40GBase-KR4, and one NC-SI (pin assignments shown in Table 2). This extension X16 Edge connector is referred as Connector B. This connector is optional.

Table 2: 1S Server Extension X16 OCP Edge Connector B Pin-Out

Default Pin-Out			
Pin Name	**B Side**	**A Side**	**Pin Name**
P12V	1	1	PRSNT_B#
P12V	2	2	P12V
P12V	3	3	P12V
GND	4	4	GND
NCSI_TXEN	5	5	NCSI_RCLK
NCSI_TXD0	6	6	NCSI_RXD0
NCSI_TXD1	7	7	NCSI_RXD1
NCSI_CRSDV	8	8	GND
NCSI_RXER	9	9	PCIE4_REFCLK_P
GND	10	10	PCIE4_REFCLK_N
PCIE3_RESET#	11	11	GND
PCIE4_RESET#	12	12	GND
PCIE5_RESET#	13	13	PCIE5_REFCLK_P
GND	14	14	PCIE5_REFCLK_N
KR4_TX0_P	15	15	GND
KR4_TX0_N	16	16	GND
GND	17	17	KR4_RX0_P
GND	18	18	KR4_RX0_N
KR4_TX1_P	19	19	GND
KR4_TX1_N	20	20	GND
GND	21	21	KR4_RX1_P
GND	22	22	KR4_RX1_N
KR4_TX2_P	23	23	GND
KR4_TX2_N	24	24	GND
GND	25	25	KR4_RX2_P

GND	26	26	KR4_RX2_N
KR4_TX3_P	27	27	GND
KR4_TX3_N	28	28	GND
GND	29	29	KR4_RX3_P
GND	30	30	KR4_RX3_N
PCIE3_REFCLK_P	31	31	GND
PCIE3_REFCLK_N	32	32	GND
GND	33	33	PCIE3_RX0_P
GND	34	34	PCIE3_RX0_N
PCIE3_TX0_P	35	35	GND
PCIE3_TX0_N	36	36	GND
GND	37	37	PCIE3_RX1_P
GND	38	38	PCIE3_RX1_N
PCIE3_TX1_P	39	39	GND
PCIE3_TX1_N	40	40	GND
GND	41	41	PCIE3_RX2_P
GND	42	42	PCIE3_RX2_N
PCIE3_TX2_P	43	43	GND
PCIE3_TX2_N	44	44	GND
GND	45	45	PCIE3_RX3_P
GND	46	46	PCIE3_RX3_N
PCIE3_TX3_P	47	47	GND
PCIE3_TX3_N	48	48	GND
GND	49	49	PCIE4_RX0_P
GND	50	50	PCIE4_RX0_N
PCIE4_TX0_P	51	51	GND
PCIE4_TX0_N	52	52	GND
GND	53	53	PCIE4_RX1_P
GND	54	54	PCIE4_RX1_N
PCIE4_TX1_P	55	55	GND
PCIE4_TX1_N	56	56	GND
GND	57	57	PCIE4_RX2_P
GND	58	58	PCIE4_RX2_N

PCIE4_TX2_P	59	59	GND
PCIE4_TX2_N	60	60	GND
GND	61	61	PCIE4_RX3_P
GND	62	62	PCIE4_RX3_N
PCIE4_TX3_P	63	63	GND
PCIE4_TX3_N	64	64	GND
GND	65	65	PCIE5_RX0_P
GND	66	66	PCIE5_RX0_N
PCIE5_TX0_P	67	67	GND
PCIE5_TX0_N	68	68	GND
GND	69	69	PCIE5_RX1_P
GND	70	70	PCIE5_RX1_N
PCIE5_TX1_P	71	71	GND
PCIE5_TX1_N	72	72	GND
GND	73	73	PCIE5_RX2_P
GND	74	74	PCIE5_RX2_N
PCIE5_TX2_P	75	75	GND
PCIE5_TX2_N	76	76	GND
GND	77	77	PCIE5_RX3_P
GND	78	78	PCIE5_RX3_N
PCIE5_TX3_P	79	79	GND
PCIE5_TX3_N	80	80	GND
GND	81	81	P12V
POWER_FAIL_N	82	82	P12V

7.4 Pin Definitions

Table 3 provides a detailed explanation of the pins. The direction of the signals is always defined from the perspective of the 1S Server module.

Table 3:Detailed Pin Definitions

Pin	Direction	Required/ Configurable	Pin Definition
P12V	Input	Required	12VAUX power from platform
I2C_SCL	Input/Output	Required	I^2C clock signal. I^2C is the primary sideband interface for server management functionality. 3.3VAUX signal. Pull-up is provided on the platform.
I2C_SDA	Input/Output	Required	I^2C data signal. I^2C is the primary sideband interface for server management functionality. 3.3VAUX signal. Pull-up is provided on the platform.
I2C_ALERT#	Output	Required	I^2C alert signal. Alerts the BMC that event has occurred that needs to be processed. 3.3VAUX signal. Pull-up is provided on the platform.
NIC_SMBUS_SCL	Input/Output	Required	Dedicated SMBus clock signal for network sideband traffic between the BMC and the NIC. 3.3VAUX signal. Pull-up is provided on the platform.
NIC_SMBUS_SDA	Input/Output	Required	Dedicated SMBus data signal for network sideband traffic between the BMC and the NIC. 3.3VAUX signal. Pull-up is provided on the platform.
NIC_SMBUS_ALERT#	Output	Required	Dedicated SMBus alert signal for network sideband traffic between the BMC and the NIC. 3.3VAUX signal. Pull-up is provided on the platform.
NCSI_RCLK	Input	Required	NC-SI reference clock for NIC
NCSI_CRSDV	Output	Required	Carrier Sense/Receive Data Valid from NIC to BMC.
NCSI_RXER	Output	Required	Receive error from NIC to BMC
NCSI_TXEN	Input	Required	Transmit enable from BMC to NIC
NCSI_RXD[0:1]	Output	Required	Receive data from NIC to BMC
NCSI_TXD[0:1]	Input	Required	Transmit data from BMC to NIC

PWR_BTN#	Input	Required	Power on signal. When driven low, it indicates that the server will begin its power-on sequence. 3.3VAUX signal. Pull-up is provided on the platform. If PWR_BTN# is held low for < 4 seconds, then this indicates a soft (graceful) power off. Otherwise, a hard shutdown is initiated.
SYS_RESET#	Input	Required	System reset signal. When driven low, it indicates that the server will begin its warm reboot process. 3.3VAUX signal. Pull-up is provided on the platform.
PRSNT_A#	Output	Required	Present signal. This is pulled low on the card to indicate that a card is installed. 3.3VAUX signal. Pull-up is provided on the platform.
PRSNT_B#	Output	Required	Extension edge connector Present signal. This is pulled low on the card to indicate that a card is installed. 3.3VAUX signal. Pull-up is provided on the platform.
COM_TX	Output	Required	Serial transmit signal. Data is sent from the 1S Server module to the BMC. 3.3VAUX signal.
COM_RX	Input	Required	Serial receive signal. Data is sent from the BMC to the 1S Server module. 3.3VAUX signal.
SVR_ID0/1/2/3 (GPIO0/1/2/3)	Input/Output	Required	Slot ID bits or open-drain GPIOs. If a system contains more than one slot, each slot will be assigned a unique ID using pull-down resistors for 0 and open for a 1. Pull-ups should be provided on the card and pull-downs should be on the platform. The server can use this slot ID to identify its location in the system. Secondly, these pins can also be use as GPIOs when they are not needed as slot IDs. 3.3VAUX signal.
KR_TX_P/N	Output	Configurable	Primary 10GBase-KR Ethernet transmit signal. Data is sent from the 1S Server module to the platform.
KR_RX_P/N	Input	Configurable	Primary 10GBase-KR Ethernet receive signal. Data is sent from the platform to the 1S Server module.
PCIE0_RESET#	Output	Required	PCIe reset signal. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.

PCIE0_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE0_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
PCIE0_REFCLK_P/_N	Output	Configurable	PCIe reference clock. This signal may or may not be connected on the platform.
PCIE1/2_RESET#	Output	Required	PCIe reset signals. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.
PCIE1_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE1_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
PCIE1_REFCLK_P/_N	Output	Configurable	PCIe reference clock. These signals may or may not be connected on the platform.
SATA0_TX_P/N	Output	Configurable	SATA 2.0 or 3.0 transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE2_RESET#	Output	Required	PCIe reset signals. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.
PCIE2_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE2_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.

PCIE2_REFCLK_P/_N	Output	Configurable	PCIe reference clocks. These signals may or may not be connected on the platform.
PCIE3_RESET#	Output	Required	PCIe reset signal. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.
PCIE3_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE3_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
PCIE3_REFCLK_P/_N	Output	Configurable	PCIe reference clocks. These signals may or may not be connected on the platform.
PCIE4_RESET#	Output	Required	PCIe reset signal. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.
PCIE4_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
PCIE4_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
PCIE4_REFCLK_P/_N	Output	Configurable	PCIe reference clock. These signals may or may not be connected on the platform.
PCIE5_RESET#	Output	Required	PCIe reset signal. If a PCIe bus is connected, this signal provides the reset signal indicating the card VRs and clocks are stable when driven high to 3.3V.
PCIE5_TX0/1/2/3_P/N	Output	Configurable	PCIe x4 bus transmit signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.

PCIE5_RX0/1/2/3_P/N	Input	Configurable	PCIe x4 bus receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
PCIE5_REFCLK_P/_N	Output	Configurable	PCIe reference clock. These signals may or may not be connected on the platform.
SATA0_RX_P/N	Input	Configurable	SATA 2.0 or 3.0 receive signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
USB_P/N	Input/Output	Required	USB 2.0 differential pair.
KR4_TX0/1/2/3_P/N	Output	Configurable	40GBase-KR4 signals. Data is sent from the 1S Server module to the platform. These signals may or may not be connected on the platform.
KR4_RX0/1/2/3_P/N	Input	Configurable	40GBase-KR4 signals. Data is sent from the platform to the 1S Server module. These signals may or may not be connected on the platform.
FAST_THROTTLE_N	Input	Required	Active low open drain signal with pull-up on 1S server. Platform generates this signal and uses it as a big hammer to throttle 1S server down to lowest possible power state as fast as possible.
POWER_FAIL_N	Input	Required	Active low open drain signal with pull-up on 1S server. When this signal is asserted by platform, it informs 1S server that base system is going to cut 12V power to 1S server in certain amount of time, which is pre-defined by base system. It is possible for 1S server to perform graceful shutdown based on this signal.

7.4.1 Required vs. Configurable Connections

The card pin-out is defined to provide basic functionality and maintain flexibility to compatible alternate/future designs. All signals marked as "Required" must be connected on the card and will be connected on the platform. All signals marked "Configurable" may be used in some applications, but not in others. They may or may not be connected on the platform or on the card. It is possible to create multiple card designs that support alternate pin assignments for the reserved signals.

All high-speed receive (RX) signals shall be placed on the A-side of the connectors and all high-speed transmit (TX) signals shall be placed on the B-side of the connector. This ensures the lowest amount of potential crosstalk between adjacent differential pairs.

7.4.2 Configurable Pin Assignment Algorithm

The platform will provide a table that defines the usage of each pair of reserved pins. This table will be accessible by over I^2C. During the power-on sequence, the SoC will query the platform and retrieve the contents of the table. If the connection type of a pair of reserved pins on the card matches the connection type of the reserved pins on the platform, those signals are enabled and the initialization sequence can begin. If a match is not found, the table cannot be retrieved. If the connections are mismatched, the SoC signals are disabled and/or powered off. The tables below provide the reserved pin mapping and encoding values for different functions.

Table 4:Byte to Configurable Pin Pair Mapping

Byte #	Byte value	Note
0	0x03	Connector A: A13/A14 PCIe0 RefClk
1	0x02	Connector A: A17/A18 PCIe0 Lane 0, Gen3
2	0x02	Connector A: A21/A22 PCIe0 Lane 1, Gen3
3	0x02	Connector A: A25/A26 PCIe0 Lane 2, Gen3
4	0x02	Connector A: A29/A30 PCIe0 Lane 3, Gen3
5	0x05	Connector A: A33/A34 SATA0, Gen3
6	0x03	Connector A: A37/A38 PCIe2 RefClk
7	0x02	Connector A: A49/A50 PCIe1 Lane 0 Gen3
8	0x02	Connector A: A53/A54 PCIe1 Lane 1 Gen3
9	0x02	Connector A: A57/A58 PCIe1 Lane 2 Gen3
10	0x02	Connector A: A61/A62 PCIe1 Lane 3 Gen3
11	0x02	Connector A: A65/A66 PCIe2 Lane 0 Gen3
12	0x02	Connector A: A69/A70 PCIe2 Lane 1 Gen3
13	0x02	Connector A: A73/A74 PCIe2 Lane 2 Gen3
14	0x02	Connector A: A77/A78 PCIe2 Lane 3 Gen3
15	0x02	Connector A: B15/B16 PCIe0 Lane 0, Gen3
16	0x02	Connector A: B19/B20 PCIe0 Lane 1, Gen3
17	0x02	Connector A: B23/B24 PCIe0 Lane 2, Gen3
18	0x02	Connector A: B27/B28 PCIe0 Lane 3, Gen3
19	0x05	Connector A: B31/B32 SATA0, Gen3
20	0x03	Connector A: B35/B36 PCIe1 RefClk
21	0x02	Connector A: B51/B52 PCIe1 Lane 0 Gen3
22	0x02	Connector A: B55/B56 PCIe1 Lane 1 Gen3
23	0x02	Connector A: B59/B60 PCIe1 Lane 2 Gen3
24	0x02	Connector A: B63/B64 PCIe1 Lane 3 Gen3
25	0x02	Connector A: B67/B68 PCIe2 Lane 0 Gen3
26	0x02	Connector A: B71/B72 PCIe2 Lane 1 Gen3

27	0x02	Connector A: B75/B76 PCIe2 Lane 2 Gen3
28	0x02	Connector A: B79/B80 PCIe2 Lane 3 Gen3
29	0x03	Connector B: A13/A14 PCIe5 RefClk
30	0x09	Connector B: A17/A18 KR4 Lane 0
31	0x09	Connector B: A21/A22 KR4 Lane 1
32	0x09	Connector B: A25/A26 KR4 Lane 2
33	0x09	Connector B: A29/A30 KR4 Lane 3
34	0x02	Connector B: A33/A34 PCIe3 Lane 0 Gen3
35	0x02	Connector B: A37/A38 PCIe3 Lane 1 Gen3
36	0x02	Connector B: A41/A42 PCIe3 Lane 2 Gen3
37	0x02	Connector B: A45/A46 PCIe3 Lane 3 Gen3
38	0x02	Connector B: A49/A50 PCIe4 Lane 0 Gen3
39	0x02	Connector B: A53/A54 PCIe4 Lane 1 Gen3
40	0x02	Connector B: A57/A58 PCIe4 Lane 2 Gen3
41	0x02	Connector B: A61/A62 PCIe4 Lane 3 Gen3
42	0x02	Connector B: A65/A66 PCIe5 Lane 0 Gen3
43	0x02	Connector B: A69/A70 PCIe5 Lane 1 Gen3
44	0x02	Connector B: A73/A74 PCIe5 Lane 2 Gen3
45	0x02	Connector B: A77/A78 PCIe5 Lane 3 Gen3
46	0x09	Connector B: B15/B16 KR4 Lane 0
47	0x09	Connector B: B19/B20 KR4 Lane 1
48	0x09	Connector B: B23/B24 KR4 Lane 2
49	0x09	Connector B: B27/B28 KR4 Lane 3
50	0x03	Connector B: B31/B32 PCIe3 RefClk
51	0x02	Connector B: B35/B36 PCIe3 Lane 0 Gen3
52	0x02	Connector B: B39/B40 PCIe3 Lane 1 Gen3
53	0x02	Connector B: B43/B44 PCIe3 Lane 2 Gen3
54	0x02	Connector B: B47/B48 PCIe3 Lane 3 Gen3
55	0x02	Connector B: B51/B52 PCIe4 Lane 0 Gen3
56	0x02	Connector B: B55/B56 PCIe4 Lane 1 Gen3
57	0x02	Connector B: B59/B60 PCIe4 Lane 2 Gen3
58	0x02	Connector B: B63/B64 PCIe4 Lane 3 Gen3
59	0x02	Connector B: B67/B68 PCIe5 Lane 0 Gen3
60	0x02	Connector B: B71/B72 PCIe5 Lane 1 Gen3
61	0x02	Connector B: B75/B76 PCIe5 Lane 2 Gen3
62	0x02	Connector B: B79/B80 PCIe5 Lane 3 Gen3

Table 5: Hex Value Encoding For Different Signal Types

Encoding	Signal Type
0x00	Off
0x01	PCIe gen 2
0x02	PCIe gen 3
0x03	PCIe clock
0x04	SATA 2.0
0x05	SATA 3.0
0x06	1000BASE-KX
0x07	10GBASE-KR
0x08	SGMII
0x09	40GBASE-KR4
0x0A-0xff	RFU

7.4.3 1S Server Edge connector Pin Assignments

The 1S server uses the default pin-out as defined in Table 1 and Table 2.

7.5 Ethernet

At least one Ethernet connection (KR_TX/RX) is required on the card. To enable maximum compatibility and a variety of potential topologies, this Ethernet port is a PHY layer device and must be capable of the following:

Table 6:Ethernet Connections

Mode	Standard	Encoding
1000BASE-X	IEEE Clause 36, 37	8b/10b
1000BASE-KX	IEEE 802.3ap	8b/10b
10GBase-KR	IEEE 802.3ap	64b/66b
40GBase-KR4	IEEE 802.3ba	64b/66b

The 1S server may use 10GBase-KR as its default Ethernet connection. Auto-negotiation and link training are required to work with the retimer or PHY on the platform.

A 40GBase-KR4 on connector B is optional.

7.5.1 Routing guidelines

To support 10Gb speeds, it is critical that the differential pairs for transmit and receive adhere to the following SoC vendor's strict route guideline.

It is strongly recommended to perform comprehensive signal integrity simulation, analysis and validation for 10GBase-KR and 40GBase-KR4 on a platform, including the 1S server, platform, connectors and vias, to determine if all the signal integrity related parameters are within the limit defined by the standards.

As a general guideline, 5dB or less channel loss budget is reserved for 10GBase-KR differential signals on the 1S server card.

7.6 SATA

The SATA connection is a SATA3.0 (6Gb/s).

7.7 USB

The USB connection is USB 2.0.

7.8 Serial port

The serial port shall be routed to the BMC on the platform. Thus, the user can access the SoC's serial console through the BMC locally or remotely via Serial-Over-LAN (SOL).

7.9 PCIe

The PCIe connection is a PCIe Gen3. A maximum of six x4 PCIe ports with dedicated PCIe reset and PCIe reference clocks are supported.

7.10 I^2C

A single I^2C connection is used to connect the BMC on the platform to the Bridge IC on the 1S server as the primary server management interface.

It shall support a minimum speed of 400kHz. However, 1MHz or higher is strongly recommended. The I2C alert signal is required and is used as an interrupt for both the Bridge IC and the BMC.

Both the BMC and the Bridge IC are I^2C master devices on the bus and they communicate with each other via the Intelligent Platform Management Bus (IPMB) protocol.

To achieve maximum bandwidth and avoid conflicts, no other devices should use this bus except for the BMC and the Bridge IC.

7.11 NIC Sideband

7.11.1 SMBus

When the SoC's integrated network controller is used as a shared NIC, its SMBus is routed to Connector A as the sideband signals. The BMC on the platform can leverage this SMBus as its out-of-band access path.

7.11.2 NC-SI

When the SoC's integrated network controller is used as a shared NIC, its NC-SI is an alternative sideband interface to the SMBus. It can run at a significant higher speed compared to the SMBus, but it requires more signals and a more complicated protocol. The BMC could use NC-SI as an alternative high-speed sideband if both the SoC and the BMC support this feature.

7.12 Slot ID and GPIO

In a multi-node platform, the 1S server often needs to know its location for server management purposes. In a BMC-free environment, slot ID bits can be used to set up unique addresses for each 1S server to identify its location.

However, due to the nature of the 1S server architecture, a BMC is always preferred on the platform side to work with the 1S server. When a BMC is present in the system, the Bridge IC shall always request its slot ID from the BMC but not probe the slot ID by itself.

There are four slot ID bits defined in the 1S server spec. These pins can be used as GPIO pins if they are not used as slot IDs.

8 Power

8.1 Input

Power for the card is provided via seven 12V pins on the primary connector and seven more 12V pins on the extension connector. Each pin supports a maximum 1.1A of current.

The nominal 12V input voltage is defined as 12.5V, +/-7%.

8.1.1 1S server power capacity

From a pure power perspective, the 1S server card is designed to support a maximum of 96W of total power consumption when only the primary connector is used, and a maximum of 192W total power consumption when both the primary and extension connectors are used.

However, the maximum power that a 1S server card can use is ultimately limited by the platform's power capacity and thermal capability. It is critical to develop a sophisticated thermal solution for the platform to keep the 1S server operating at a safe condition when it runs with a maximum power load.

8.1.2 Power sequence and standby power

Because there is only one 12V power input to the 1S server, there is not a power sequence requirement to power on the 1S server card from the platform perspective. The 1S server card itself usually has its own power sequencing requirement.

However, a standby 3.3V_AUX power rail on the card is required to power the Bridge IC at all power states. The SoC and its control circuits may require some stand-by power rails. For example, when the BMC uses the SoC's integrated network controller as a shared NIC, that NIC requires stand-by power rails to keep the out-of-band function alive when the SoC is in standby mode.

It is the designer's responsibility to provide proper standby power rails from the main 12V_AUX input with possible specific power sequencing. Care must be taken to avoid any possible leakage path between the power domains, such as 1S server stand-by, 1S server payload, and platform power.

8.1.3 Standby power budget

The 1S server shall consume less than 5W when it operates in standby mode.

8.2 Hot-Plug Support

Hot-plug is not supported by the 1S server specification.

8.3 VR Efficiency

All Voltage Regulators (VRs) providing over 15W on the card are at least 91% efficient when loaded between 30% and 90% of the full load.

8.4 Input Capacitance

The capacitance on the input 12V rail of the 1S server is recommended not to exceed 1000uF. However, platform designers must make certain that the overall capacitance on the 12.5V rail of the entire platform meets the system's power supply requirement and does not cause instability.

8.5 Power reading and power capping

The 1S server shall implement sophisticated power management features.

The 1S server shall have the power monitoring capability to read power consumption reliably and accurately. Preferably, the 1S server needs to provide a one-second average power reading with 3% accuracy. This power consumption number should be accessible from in-band and out-of-band. As shown in figure 3-1, a power sensor with an I^2C interface shall be used at the 12.5V input side so that either the Bridge IC or SoC can read the whole-card power consumption and send this information to the BMC per its request.

The 1S server shall have the ability to perform power capping as per the platform BMC's request. As a minimum requirement, the 1S server shall be able to throttle itself down to lowest possible power state as quickly as possible when the platform asserts the FAST_THROTTLE_N signal or sends such a request to the Bridge IC and SoC by BMC.

A preferred power-capping implementation is to reduce the 1S server's power consumption gradually with fine-grained power control by steps as small as 5 watts and to reach the control target power limit within 3 seconds. This process shall be smooth but fast, and the settled power value shall be within -3% of the target power limit set by the platform.

The platform can generate a POWER_FAIL_N signal to inform the 1S server that the 12V input power to the server is going to be cut off in certain amount of time (which is pre-defined by the platform). The 1S server can leverage this signal to develop mechanisms to protect critical data prior to a power outage.

9 Functional

9.1 System on a Chip

The 1S server shall use a single socket SoC that incorporates a multi-core CPU, memory controller, and other IO devices.

9.2 Memory

The 1S server shall support DDR4 memory.

9.3 NIC

The optional integrated NIC in the SoC shall support 10GBASE-KR, IPv4, IPv6, and iSCSI boot. Having iSCSI and TCP/IP offload capabilities are highly desirable.

9.4 Storage

The 1S server supports a minimum 128GB of high-performance local storage, such as SSD. The M.2 form factor is preferred for the boot device.

SATA-based storage device must be connected to SATA port 0 of the SoC. Additional SATA connections are connected to port 1 and higher. It is recommended to route a PCIe link to the M.2 connector as a BOM option that a PCIe based NVMe solid-state drive can be supported for applications that desire high disk performance.

9.5 EEPROM

The 1S server includes an I^2C-accessible Electrically Erasable Programmable Read-Only Memory (EEPROM). The EEPROM must be accessible from the platform via the Bridge IC. Its capacity must be at least 128Kbits. The EEPROM will contain the Field Replaceable Unit Identification (FRU ID) information and any additional configuration information that may be required. The FRU ID is formatted in accordance with the IPMI Platform Management FRU Information Storage Definition document.

The EEPROM must contain the following:

- Board Manufacturer
- Board Name
- Board Serial Number
- Board Part Number
- Product Manufacturer
- Product Name
- Product Part Number
- Product Serial Number
- Product Asset Tag
- Product Build: e.g. EVT, DVT, PVT, MP
- Product Version: e.g. C1
- Manufacturing Date and Time
- Manufacturing Lot Code: (preferred, but optional)
- Manufacturing Work Order: (preferred, but optional)

- PCB Revision
- SoC Model Name/Number
- SoC Revision
- SoC Tj_{MAX} (Maximum Junction Temperature)

This information will be available via IPMI commands from the BMC.

9.6 BIOS

The card supplier is responsible for supplying and customizing the BIOS for the SoC. The requirements are outlined in this section.

9.6.1 UEFI

The BIOS shall be a UEFI compatible BIOS.

9.6.2 Configuration and Features

The BIOS is tuned to minimize card power consumption. It has the following features:

- Disables unused devices, including PCIe lanes, USB ports, SATA/SAS ports, etc.
- A BIOS setup menu
- The SoC settings can be tuned to achieve the optimal combination of performance and power consumption.

9.6.3 BIOS Settings Tools

The card supplier shall provide a tool to make BIOS setting changes without requiring a BIOS re-flash. The BIOS settings update tool must also support success and failure codes so that updates can be easily scripted.

9.6.4 PXE Boot

The BIOS supports PXE boot and provides the ability to modify the boot sequence. When PXE booting, the card first attempts to boot from the first available Ethernet device .

The default boot device priority is:

1. Network (search all configured network interfaces)
2. HDD, SSD, or flash device (local or remote)
3. CD-ROM
4. Removable Device

This process loops indefinitely and requires no user intervention.

9.6.5 iSCSI Boot

The BIOS shall be capable of iSCSI network boot.

9.6.6 Other Boot Options

The BIOS also supports booting from SATA/SAS and USB interfaces. The BIOS provides the capability to select different boot options.

9.6.7 BIOS Update

The BIOS can be updated from the OS under these scenarios:

- Scenario 1: Sample/Audit BIOS settings
    - Return current BIOS settings, or
    - Save/export BIOS settings in a human-readable form that can be restored/imported (as in Scenario 2)
- Scenario 2: Update BIOS with pre-configured set of BIOS settings
    - Update/change multiple BIOS settings
    - Reboot
- Scenario 3: BIOS/firmware update with a new revision
    - Load new BIOS/firmware on machine and update, retain the current BIOS settings
    - Reboot

Additionally, the update tools have the following capabilities:

- Update from the operating system.
- Can complete BIOS update or setup change with a single reboot (no PXE boot, no multiple reboots)
- No user interaction (such as prompts)
- BIOS updates and option changes do not take longer than five minutes to complete

9.6.8 Remote BIOS Update

The BIOS can be updated remotely under these scenarios:

- Scenario 1: Sample/Audit BIOS settings
    - Return current BIOS settings, or
    - Save/export BIOS settings in a human-readable form that can be restored/imported (as in Scenario 2)
- Scenario 2: Update BIOS with a pre-configured set of BIOS settings
    - Update/change multiple BIOS settings
    - Reboot
- Scenario 3: BIOS/firmware update with a new revision
    - Load new BIOS/firmware on machine and update, retain the current BIOS settings
    - Reboot

Additionally, the update tools have the following capabilities:

- Update from the remote host over the LAN connection to BMC
- Can complete BIOS update or setup change with a single reboot (no PXE boot, no multiple reboots)
- No user interaction (such as prompts)
- BIOS updates and option changes do not take longer than 20 minutes to complete
- Can be scripted and propagated to multiple machines

9.6.9 SMBIOS Event Log

Per the SMBIOS specification Version 2.6, the BIOS implements SMBIOS Type 15 for an event log; the assigned area is large enough to hold more than 500 event records (assuming the maximum event record length is 24 bytes, then the size will be larger than 12KB), and follows the SMBIOS event log organization format for the event log.

A system access interface and application software must be provided to retrieve and clear the event log from the BIOS, including, at minimum, a Linux application for the CentOS operating system and driver as needed. The event log must be retrieved and stored as a readable text file that is easy to handle by a scripting language under Linux. Each event record includes enhanced information identifying the error source device's vendor ID, card slot ID, and device ID.

9.6.10 Logged Errors

The following list of errors are logged by the BIOS or the Bridge IC. These errors must include the date, time, and location information so that failing components can be easily identified.

- CPU/Memory errors: Both correctable ECC and uncorrectable ECC errors are logged into the event log. Error categories include DRAM, Link, and L3 cache.
- PCIe* errors: Any errors that have a status register are logged into the event log, including root complex, endpoint device, and any switch upstream/downstream ports if available. Link disable on errors are also be logged. Fatal, non-fatal, or correctable error classification follows the chipset vendor's recommendation.
- POST errors: All POST errors detected by the BIOS during POST are logged into the event log.
- SATA or SAS errors: All correctable and uncorrectable errors are logged.
- System reboot events
- Sensor values exceeding warning or critical thresholds

9.6.11 Error Thresholds

An error threshold setting must be enabled for both correctable and uncorrectable errors. Once the programmed threshold is reached, an event is triggered and logged.

- Memory Correctable ECC: The threshold default value is 1,000. When the threshold is reached, the BIOS logs the event and includes the physical DIMM location.

9.6.12 POST Codes

The BIOS outputs a set of power-on Self-Test (POST) codes identifying the current initialization step and any errors encountered along the initialization. The output is provided on the serial console and errors are logged.

During the boot sequence the BIOS shall initialize and test each DIMM module. If a module fails initialization or does not pass the BIOS test, the following POST Codes should indicate which DIMM has failed:

- The first hex character indicates which CPU interfaces the DIMM module
- The second hex character indicates the number of the DIMM module.
- The POST Code will also display both the error major code and minor code from the memory reference code.

- The display sequence will be “00”, DIMM location, Major code and Minor code with a one second delay for every code displayed.
- The BIOS shall repeat the display sequence indefinitely to allow time for a technician to service the system.

A DIMM location code table is provided in Table 7.
DIMM number count starts from the furthest DIMM from the SoC.

Table 7: DIMM Error Code Table

Code	Result
A0	Channel 0 DIMM 0 Failure
A1	Channel 0 DIMM 1 Failure
B0	Channel 1 DIMM 0 Failure
B1	Channel 1 DIMM 1 Failure

9.7 1S Server Management

The primary server management functions will be provided using a BMC on the platform. The BMC on the platform will use an I^2C bus as the management interface. This section identifies the required information that must be accessible from the BMC.

9.7.1 Bridge IC

This 1S server design specification is SoC agnostic. To accommodate different types of SoCs, a common Bridge IC is defined as the bridging device between the SoC and the BMC. The Bridge IC shall be on stand-by power so that it can be accessed by the BMC even when the SoC is powered down.

On the platform side, the BMC and Bridge IC communicate with each other with IPMI messages over the I^2C bus. To enable prompt communication, this I^2C bus shall be a point-to-point link without any other devices on the same bus. It shall run in high-speed mode with a minimum speed of 400KHz. When possible, a 1MHz or better speed is strongly recommended.

On the 1S server side, the Bridge IC shall have FRU EEPROM and thermal sensors on a local I^2C bus. The BMC shall be able to communicate with the Bridge IC to inquire the FRU and thermal data through IPMB. The FRU EEPROM’s data format is defined in Section 8.5. The thermal sensors are mainly used to measure the inlet and outlet temperatures of the 1S server for the platform thermal management’s algorithm.

If the SoC has an internal management microcontroller that supports IPMB, the Bridge IC shall connect this microcontroller to one of its I^2C port. In this way, the Bridge IC now behaves as a transparent bridge to forward IPMI messages between the BMC and the SoC’s internal microcontroller. With this transparent bridge, the BMC can directly work with the SoC’s own management controller to perform most server management functions.

The Bridge IC shall also implement a system interface, such as KCS or SSIF, to enable in-band server manageability. If KCS is implemented, it shall support both standard SMM and SMS interfaces.

The Bridge IC shall monitor the 1S server's sensors, such as voltage sensors, power sensors, and digital sensors for critical GPIOs. The BMC shall be able to inquire about the 1S server's status by reading these sensors and taking actions via the Bridge IC.

It is recommended to use a versatile microcontroller as the Bridge IC. The microcontroller has a compact size, uses a low amount of power, and has adequate functions to support all required bridging functions. This microcontroller needs to have two Serial Peripheral Interfaces (SPI) interfaces. The first SPI bus is used as its own boot ROM. The second SPI can be used to re-program the SoC's boot ROM when it is corrupted with a multiplexer and BMC's support.

9.7.2 I^2C addressing

The 1S server and BMC will communicate using IPMI 2.0 commands transmitted over the I^2C connection through a Bridge IC on the 1S Server card. The I^2C bus address for the Bridge IC shall be configured as 0x40. The BMC on the platform is configured as 0x20.

9.7.3 Message Transfer

As the bridge between the server SoC and BMC on the platform, the Bridge IC shall provide ways to transfer messages between them via various interfaces like KCS, SMBus, UART, and I^2C.

For in-band management, the Bridge IC shall be able to forward the SoC's Keyboard Controller Style (KCS) request to the BMC and then send the received response back to the SoC.

When the BMC sends a request on the I^2C bus meant for the SoC-vendor specific management controller, the Bridge IC shall forward the command on I^2C bus and send received response back to the BMC.

It is possible to implement an alternative SOL through the Bridge IC. When the alternative SOL feature is enabled, the serial data from the SoC's serial port shall be sent to the BMC via the I^2C. When the BMC sends SOL data, it shall be emitted via the serial port.

9.7.4 Platform Discovery and Configuration

The Bridge IC shall provide a way for the BMC to discover platform capabilities such as electrical interface assignment. It shall provide a way to discover and configure its own capabilities like enabling or disabling the SOL interface and/or POST code interface.

9.7.5 POST Code Access

During the power-on stage, the SoC usually sends out status/error information on the POST Code interface. The Bridge IC needs to provide a way for BMC to access POST Code information. The Bridge IC shall keep the latest POST Code in a 230-byte buffer. Whenever the BMC is available, the Bridge shall send the POST Code as soon as it is received on the POST Code interface. Whenever the BMC is not available (such as the BMC being in a firmware update mode or during BMC boot-up), the Bridge IC shall add the latest POST Code at the top of the 230-byte buffer. The BMC shall be able to retrieve the POST Code buffer from Bridge with the latest POST Code.

9.7.6 IPMB Interface

The Bridge IC shall provide an IPMB interface for the BMC to access various IPMI resources on the 1S Server. To meet this requirement, the Bridge IC shall implement various standard IPMI commands. It shall implement FRUID commands to identify the 1S server, System Event Log

(SEL) commands to store 1S Server specific event logs, and Sensor Data Repository (SDR) commands to identify various sensors described for the specific 1S server.

9.7.7 Firmware Update

The Bridge IC shall be able to update the firmware of the programmable devices on the 1S Server, such as boot firmware (BIOS, UBoot, UEFI and so on), the SoC-vendor's specific management controller firmware, CPLD/FPGA image, and the Bridge IC's firmware.

The Bridge shall provide a way for the BMC to access the version information of various firmware components on the 1S Server, initiate the update process for various firmware components on the 1S Server, and detect and retransmit corrupted firmware image packets during transit from the BMC to the Bridge IC.

9.7.8 Network status LED Control

When the SoC's integrated network controller is used, often the network status LEDs are physically located on platform side. This is not the case on the 1S server. In this case, the Bridge IC needs to collect various network controller statuses, such as network speed, link status, and activity. The Bridge IC sends this information to the BMC through the I^2C bus between them. The BMC then can then use the LED information to control LEDs on the platform side. The Bridge IC shall provide a way for the BMC to retrieve the required network controller's LED status.

9.7.9 GPIO Register

The Bridge IC shall provide a GPIO interface to the BMC through the GPIO register block. In this way, the BMC can control the GPIO behind the Bridge IC (or this hardware abstraction layer) by accessing this register block.

The Bridge IC shall provide a way for the BMC to configure the GPIO pin's direction and interrupt capability, and provide a way to get/set the current status of the GPIO signals. The Bridge IC shall send an interrupt message to the BMC when the interrupt enabled GPIO signal changes it's state. The GPIO register interface exposed by the Bridge IC shall provide four bytes to represent 32 signals that indicate various conditions as shown in Table 8.

Table 8: Bridge IC GPIO Table

No.	GPIO offset	GPIO Pin Function	Comments
1	Byte 1 -bit [0]	Power Good - CPU	Indicates CPU's Power input is good
2	Byte 1 - bit [1]	Power Good – AUX Comp.	Indicates AUX component's Power input is good
3	Byte 1 - bit [2]	VR HOT – DDR Voltage	Indicates DDR Voltage Regulator is hot
4	Byte 1 - bit [3]	VR HOT – CPU Core Voltage	Indicates CPU Core Voltage Regulator is hot
5	Byte 1 - bit [4]	CPU Throttle	Indicates CPU is throttled to lowest power
6	Byte 1 - bit [5]	AUX Comp. Throttle	Indicates AUX Component is throttled
7	Byte 1 - bit [6]	Memory Over Temp	Indicates one or more DIMM(s) has experienced Over Temperature condition

8	Byte 1 - bit [7]	CPU Over Temp	Indicates CPU is experiencing Over Temperature condition
9	Byte 2 - bit [0]	AUX Comp. Over Temp	Indicates AUX. Component is experiencing Over Temperature condition
10	Byte 2 - bit [1]	CPU Internal VR Error	Indicates CPU internal Voltage Regulator Error condition
11	Byte 2 - bit [2]	CPU Catastrophic Error	Indicates CPU experienced Catastrophic Error
12	Byte 2 - bit [3]	CPU Non-Recoverable Error	Indicates CPU experienced Non Recoverable Error
13	Byte 2 - bit [4]	CPU Critical Error	Indicates CPU experienced Critical Error
14	Byte 2 - bit [5]	CPU Non-Critical Error	Indicates CPU experienced Non-Critical Error
15	Byte 2 - bit [6]	S4 Sleep State Transition	Indicates CPU entered S4 sleep state
16	Byte 2 - bit [7]	NMI	Non-Maskable Interrupt
17	Byte 3 - bit [0]	SMI	System Management Interrupt
18	Byte 3 - bit [1]	CPU POR Reset	Initiate CPU Power On Reset
19	Byte 3 - bit [2]	FP Reset Button	Initiate Reset to the CPU resulted from Front Panel Reset button press action
20	Byte 3 - bit [3]	CPU Reset	Initiate CPU Reset
21	Byte 3 - bit [4]	POST Completed	Indicates that POST phase has completed
22	Byte 3 - bit [5]	S3 Sleep State Transition	Indicates that CPU has entered S3 sleep state
23	Byte 3 - bit [6]	Reserved	
24	Byte 3 - bit [7]	Power Good – CPU Core Voltage	Indicates that CPU Core Voltage Power is Good
25	Byte 4 - bit [0]	Slot ID0	Slot ID0 – Slot ID3 provides location information
26	Byte 4 - bit [1]	Slot ID1	
27	Byte 4 - bit [2]	Slot ID2	
28	Byte 4 - bit [3]	Slot ID3	
29	Byte 4 - bit [4]	BMC Ready	Used by the BMC after initialization to inform the CPU that it is ready to accept commands
30	Byte 4 - bit [5]	Reserved	
31	Byte 4 - bit [6]	Reserved	
32	Byte 4 - bit [7]	Reserved	

9.7.10 IPMI commands

The Bridge IC must support the IPMI commands shown in Table 9.

Table 9: Bridge IC supported IPMI command Table

IPMI Command	Net Function	CMD#
Get Device ID	App	01h
Get Self Test Results	App	04h
Get System GUID	App	37h
Master Write-Read I^2C	App	52h
Get FRU Inventory Area Info	Storage	10h
Read FRU Inventory Data	Storage	11h
Write FRU Inventory Data	Storage	12h
Get SDR Repository Info	Storage	20h
Reserve SDR Repository	Storage	22h
Get SDR	Storage	23h
Get SEL Info	Storage	40h
Get SEL Allocation Info	Storage	41h
Reserve SEL	Storage	42h
Get SEL Entry	Storage	43h
Add SEL Entry	Storage	44h
Clear SEL	Storage	47h
Get Sensor Reading	Sensor/Event	2Dh
Send request message to BMC	OEM (0x38)	01h
Send request message to Bridge-IC	OEM (0x38)	02h
Get all GPIO status	OEM (0x38)	03h
Set all GPIO status	OEM (0x38)	04h
Get GPIO configuration	OEM (0x38)	05h
Set GPIO configuration	OEM (0x38)	06h
Send interrupt to BMC	OEM (0x38)	07h
Send POST Code to BMC	OEM (0x38)	08h
Request POST Code data	OEM (0x38)	12h
Firmware Update	OEM (0x38)	09h
Firmware Verify	OEM (0x38)	0Ah
Get Firmware version	OEM (0x38)	0Bh

Enable Bridge IC update flag	OEM (0x38)	0Ch
Get NIC LED Frequency	OEM (0x38)	0Dh
Bridge IC Discovery	OEM (0x38)	0Eh
Platform Discovery	OEM (0x38)	0Fh
Set Bridge IC Configuration	OEM (0x38)	10h
Bridge IC Reset Cause	OEM (0x38)	11h
Bridge IC enter update ode	OEM (0x38)	13h
Set System GUID	OEM (0x38)	EFh

Table 10 provides details of the IPMI Original Equipment Manufacturer (OEM) commands that are defined in Table 9.

Table 10: Bridge IC supported IPMI command Table

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
01h	Send request message to BMC	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Request interface 01h: Management Controller 02h: SOL 03h: KCS Byte 5:X – Request data Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2 – Request interface 01h: Management Controller 02h: SOL 03h: KCS Byte 3:X – Response data	This command is used for the Bridge IC to transfer requests to the BMC. For example: 1. The Bridge IC gets "Get Device ID" command 0x06 0x01 from KCS. 2. The Bridge IC will send this command to the BMC as below. 0x38 0x01 0x03 0x06 0x01 3. The BMC responds Get Device ID data.
02h	Send request message to Bridge IC	Request: 4.1 Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 –Receive interface 01h: Management Controller 02h: SOL Byte 5:X – Request data from BMC Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2 –Receive interface 01h: Management Controller 02h: SOL Byte 3:X – Response data	This command is used for the BMC to send requests to the Bridge IC. For example: 1. When the BMC wants to send "Get Device ID" command to the management controller it can use this command: 0x38 0x02 0x01 0x06 0x01 2. When the Bridge IC receives this command, it will send "Get Device ID" command to the management controller and get the response from the management controller. 3. The Bridge IC responds with this command to the BMC.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
03h	Get all GPIO status	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:5 – Get all GPIO status 0b: Low 1b: High	This command used by the BMC to get GPIO status from the Bridge IC. Refer to Table 8 GPIO mapping table.
04h	Set all GPIO status	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4:7 – GPIO enable mask, refer to GPIO mapping table 0b: Disable 1b: Enable Byte 8:11 – Set all GPIO status 0b: Low 1b: High Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	This command used by the BMC to set GPIO status from the Bridge IC. Refer to Table 8 GPIO Mapping Table.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
05h	Get GPIO configuration	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4:7 – GPIO enable mask, refer to GPIO mapping table 0b: Disable 1b: Enable Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:X – GPIO configuration (one byte for one GPIO pin configuration) Bit[0] – Input/output pin 0b : Input pin 1b : Output pin Bit[1] – interrupt disable/enable 0b : Disable 1b : Enable Bit[2] – Edge trigger 0b : Edge trigger (default) Bit[3:4] – Trigger type 00b : Falling edge 01b : Rising edge 10b : Both 11b : Reserved	This command used by the BMC to get the GPIO configuration from the Bridge IC. Refer to Table 8 GPIO Mapping Table.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
06h	Set GPIO configuration	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4:7 – GPIO enable mask, refer to GPIO mapping table 0b: Disable 1b: Enable Byte 8:X – GPIO configuration (one byte for one GPIO pin configuration) Bit[0] – Input/output pin 0b : Input pin 1b : Output pin Bit[1] – interrupt disable/enable 0b : Disable 1b : Enable Bit[2] – Edge trigger 0b : Edge trigger (default) Bit[3:4] – Trigger type 00b : Falling edge 01b : Rising edge 10b : Both 11b : Reserved Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	This command used by the BMC to set the GPIO configuration to the Bridge IC. Refer to Table 8 GPIO Mapping Table.
07h	Send interrupt to BMC	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Interrupt GPIO number, refer to GPIO mapping table Byte 5 – Trigger type 00h: Falling edge 01h: Rising edge Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	This command is used for Interrupt notification from the Bridge IC to the BMC. Refer to Table 8 GPIO Mapping Table.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
08h	Send POST Code to BMC	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Data length Byte 5:X – POST Code data Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	The Bridge IC supports a maximum of 230 bytes to buffer BIOS POST Code when the BMC is not ready. The POST Code data will be in FIFO manner i.e. with the first POST Code as the first byte. In case the BMC is ready, the Bridge IC will send one POST Code to BMC at a time.
12h	Request POST Code data	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:X – POST Code data	The BMC can get all POST Code data by this command. The Bridge IC will buffer POST Code data for last boot. The POST Code data will be in LIFO manner i.e. with the latest POST code as the first byte. The Bridge IC clears the buffer when the system powers on. The maximum buffer data length is 230 bytes.
09h	Firmware Update	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – update target 00h: BIOS 01h: CPLD Byte 5:8 – Offset Byte 9:10 – Data length Byte11:X – Update image data Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) 80h –Write flash error 81h – Power status check fail 82h – Data length error 83h – Flash erase error	This command is used to update the BIOS and CPLD Firmware from the BMC.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
0Ah	Firmware Verify	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 –update target 00h: BIOS 01h: CPLD Byte 5:8 – Offset Byte 9:10 – Data length Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) 80h – Checksum error 82h – Data length error 84h – Read flash error Byte 2:5 – Checksum	This command is used to verify the BIOS and CPLD Firmware from the BMC.
0Bh	Get Firmware version	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 –update target 00h: BIOS 01h: CPLD 02h: Bridge IC 03h: Management Controller firmware version Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:X – BIOS version - 8 bytes length ASCII data, ex: F07_1A01 CPLD version - TBD Bridge IC version – 2 bytes length, ex: 1.03. Return data will be 0x01 0x03. Management Controller's firmware version – 5 bytes length, ex: version 03.0.0.010. Return data will be 0x03 0x00 0x00 0x01 0x00.	

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
0Ch	Enable Bridge IC update flag	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Enable update interface flag 00h: UART 01h: I2C 02h: LPC Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	This command is used to enable the Bridge IC update flag from the BMC.
0Dh	Get NIC LED frequency	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2 – LED frequency 00h: No blinking 01h: Solid on 02h: Slow flashing 03h: Fast flashing	
0Eh	Bridge IC discovery	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2 – Bit[0] – SOL interface 0b : Disable 1b : Enable Bit[1] – POST Code 0b : Disable, Bridge IC will not send post code to BMC 1b : Enable	

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
0Fh	Platform discovery	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2 – 0x03: A13/A14 PCIe0 RefClk Byte 3 – 0x02: A17/A18 PCIe0 Lane 0, Gen3 Byte 4 – 0x02: A21/A22 PCIe0 Lane 1, Gen3 Byte 5 – 0x02: A25/A26 PCIe0 Lane 2, Gen3 Byte 6 – 0x02: A29/A30 PCIe0 Lane 3, Gen3 Byte 7 – 0x05: A33/A34 SATA0, Gen3 Byte 8 – 0x03: A37/A38 PCIe2 RefClk Byte 9 – 0x02: A49/A50 PCIe1 Lane 0 Gen3 Byte 10 – 0x02: A53/A54 PCIe1 Lane 1 Gen3 Byte 11 – 0x02: A57/A58 PCIe1 Lane 2 Gen3 Byte 12 – 0x02: A61/A62 PCIe1 Lane 3 Gen3 Byte 13 – 0x02: A65/A66 PCIe2 Lane 0 Gen3 Byte 14 – 0x02: A69/A70 PCIe2 Lane 1 Gen3 Byte 15 – 0x02: A73/A74 PCIe2 Lane 2 Gen3 Byte 16 – 0x02: A77/A78 PCIe2 Lane 3 Gen3 Byte 17 – 0x02: B15/B16 PCIe0 Lane 0, Gen3 Byte 18 – 0x02: B19/B20 PCIe0 Lane 1, Gen3 Byte 19 – 0x02: B23/B24 PCIe0 Lane 2, Gen3 Byte 20 – 0x02: B27/B28 PCIe0 Lane 3, Gen3 Byte 21 – 0x05: B31/B32 SATA0, Gen3 Byte 22 – 0x03: B35/B36 PCIe1 RefClk Byte 23 – 0x02: B51/B52 PCIe1 Lane 0 Gen3 Byte 24 – 0x02: B55/B56 PCIe1 Lane 1 Gen3 Byte 25 – 0x02: B59/B60 PCIe1 Lane 2 Gen3 Byte 26 – 0x02: B63/B64 PCIe1 Lane 3 Gen3 Byte 27 – 0x02: B67/B68 PCIe2 Lane 0 Gen3 Byte 28 – 0x02: B71/B72 PCIe2 Lane 1 Gen3 Byte 29 – 0x02: B75/B76 PCIe2 Lane 2 Gen3 Byte 30 – 0x02: B79/B80 PCIe2 Lane 3 Gen3	For scalability needs, we list all configurations in this command response. If we store that information in the BMC FW and check by FRU ID, we will need to modify the BMC FW whenever there is a new card. We recommend storing the info on MB.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
10h	Set Bridge IC configuration	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Config Bridge IC Bit[0] – SOL interface 0b : Disable 1b : Enable Bit[1] – POST Code 0b : Disable, Bridge IC will not send post code to BMC 1b : Enable Bit[2] – KCS interface 0b : Disable, Bridge IC will not send KCS command to BMC 1b : Enable Bit[1] – IPMB message 0b : Disable, Bridge IC will not send IPMB message to BMC 1b : Enable Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	When the BMC enters update mode, the BMC can use this command to disable various communication to the Bridge IC.
11h	Bridge IC reset cause	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Reset cause 0x00 : Cold reset by Firmware update 0x01 : Watchdog timeout Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	The Bridge IC will send this command to notify the BMC when the Bridge IC is reset.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
13h	Bridge IC enter update mode	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 2 – Firmware mode 0x01 : normal mode 0x0F : update mode Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	This command is used to notify the BMC. The Bridge IC enters update mode and normal mode before and after the Firmware updates.
EFh	Set System GUID	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4:19 – System GUID. See Picture 5, GUID format Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)	

9.7.11 Management interface

The SoC could have its own management controller. However, this controller has to work together with the Bridge IC as well as the BMC to perform server management tasks. The management interface between the Bridge IC and the SoC consists of a supplier-agnostic interface that combines a simple register interface to abstract the card-specific details. This interface can be implemented in hardware, software, or a combination of the two.

9.7.12 Serial Console

The SoC provides a serial UART that is connected directly to the card edge. This connection will be used as the BIOS or OS serial console and will also be available as an SOL connection via the BMC. The BIOS menus must be fully accessible and text-based. Any hot keys that are required must be transmittable through a serial console session.

The BIOS should default to 57,600 bps/8N1.

9.7.13 Power Control

The BMC controls power on, off, and reset directly via the signals defined in the pin-out. If 12V to the card is lost and returns (“AC Lost”), the BIOS must be configurable to enable either an immediate or delayed power-on, or the last power state prior to the event.

9.7.14 Thermal Alerts

The SoC provides a mechanism to provide thermal alerts and over temperature notifications. The BMC must be able to receive these alerts in a timely fashion to allow it take action quickly. The `I2C alert` signal must be used. In some cases, an over temperature condition may occur which forces the SoC to power-off immediately. This condition must be logged.

9.7.15 Sensors

The following list of analog and discrete sensors are provided and are reported by the Bridge IC to the BMC.

Analog sensors include:

- Outlet Temperature
- Inlet Temperature
- VR Temperature(s)
- VR Current(s)
- SoC Temperature
- SoC Thermal Margin
- DIMM Temperature(s)
- SoC Package Power
- SoC Tj_{MAX}
- Voltage Sensor(s)
- Current Sensor(s)
- Power(s)

Discrete sensors include:

- CPU Thermal Trip
- System Status
- SoC Fail
- System Boot Status
- SoC/DIMM Hot
- VR Hot

Event Only Sensors Include:

- Power Threshold Event
- POST Error
- Power Error
- SoC Throttle
- Machine Check Error
- PCIe Error
- Other IIO Error
- Memory ECC Error

9.8 LEDs

Each card includes a power LED that illuminates when the card's power-on sequence has completed successfully. The LED is blue in color and placed on the leading edge of the card (cold-aisle). Flashing this LED may also be used to identify a card.

10 Environmental Requirements

The full system with the server card installed meets the following environmental requirements:

- Gaseous contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Storage relative humidity: 10% to 90% (non-condensing)
- Storage temperature range: -40°C to +70°C (long-term storage)
- Transportation temperature range: -55°C to +85°C (short-term storage)
- Operating altitude with no de-rating to 1,000 meters (3,300 feet)

10.1 Vibration and Shock

The motherboard meets shock and vibration requirements according to the following IEC specifications: IEC78-2-(*) and IEC721-3-(*) Standard & Levels.

Table 11: Vibration and Shock Requirements

	Operating	Non-Operating
Vibration	0.5g acceleration, 1.5mm amplitude, 5 to 500 Hz, 10 sweeps at 1 octave / minute for each of the three axes (one sweep is 5 to 500 to 5 Hz)	1g acceleration, 3mm amplitude, 5 to 500 Hz, 10 sweeps at 1 octave / minute for each of the three axes (one sweep is 5 to 500 to 5 Hz)
Shock	6g, half-sine 11ms, 5 shocks for each of the three axes	12g, half-sine 11ms, 10 shocks for each of the three axes

11 Prescribed Materials

11.1 Disallowed Components

The following components are not used in the design of the motherboard:

- Components disallowed by the European Union's Restriction of Hazardous Substances Directive (RoHS 6)
- Trimmers and/or potentiometers
- Dip switches

11.2 Capacitors and Inductors

The following limitations apply to the use of capacitors:

- Only aluminum organic polymer capacitors made by high-quality manufacturers are used and must be rated at 105°C.
- All capacitors have a predicted life of at least 50,000 hours at 45°C inlet air temperature, under the worst conditions.
- Tantalum capacitors using manganese dioxide cathodes are forbidden.
- Surface Mount (SMT) ceramic capacitors with a case size greater than 1206 are forbidden. The 1206 case size is still allowed when installed far from the PCB edge and with a correct orientation that minimizes the risk of cracks.
- Ceramic material for SMT capacitors must be X7R or better (COG or NP0 type are used in critical portions of the design). Only SMT inductors may be used. The use of through-hole inductors is disallowed.

11.3 Component De-rating

For inductors, capacitors, and FETs, de-rating analysis is based on at least 20% de-rating.

12 Labels and Markings

The motherboard shall include the following labels on the component side of the motherboard. The labels shall not be placed in a way, which may cause them to disrupt the functionality or the airflow path of the motherboard.

Description	Type	Barcode Required?
Safety Markings	Silk Screen	No
Vendor P/N, S/N, REV (Revision would increment for any approved changes)	Adhesive label	Yes
Vendor Logo, Name & Country of Origin	Silk Screen	No
PCB Vendor Logo, Name	Silk Screen	No
Date Code (Industry Standard: WEEK / YEAR)	Adhesive label	Yes
RoHS compliance	Silk Screen	No
WEEE Symbol. The motherboard will have the crossed out wheeled bin symbol to indicate that it will be taken back by the Manufacturer at the end of its useful life. This is defined in the European Union Directive 2002/96/EC of January 27, 2003 on Waste Electrical and Electronic Equipment (WEEE) and any subsequent amendments.	Silk Screen	No
CE Marking	Silk Screen	No
UL Marking	Silk Screen	No

13 Revision History

Author	Description	Revision	Date
Yan Zhao	▪ Initial draft.	0.1	12/20/2014
Yan Zhao	▪ Incorporated review comments.	0.2	02/09/2015
Yan Zhao	▪ Incorporated review comments.	0.3	02/16/2015
Yan Zhao	▪ Incorporated review comments.	0.4	01/12/2016