

OPEN

Compute Project

Twin Lakes 1S Server Design Specification

V1.00

Author:

Jon Ehlen, Mechanical Engineer, Facebook

Damien Chong, Hardware Engineer, Facebook

Ben Wei, Software Engineer, Facebook

Jarrold Clow, Thermal Engineer, Facebook

Copyrights and Trademarks

Intel® is a trademark of Intel Corporation in the U.S. and/or other countries.

Xeon® is a trademark of Intel Corporation in the U.S. and/or other countries.

Texas Instruments™ is a trademark of Texas Instruments Incorporated.

Tiva™ is a trademark of Texas Instruments Incorporated.

1 Scope

This specification describes the design of the Twin Lakes 1S server based on the Intel® Xeon® Processor D-2191 System-on-a-Chip (SoC).

2 Contents

Copyrights and Trademarks.....	2
1 Scope.....	3
2 Contents	3
3 Overview	5
4 License.....	8
5 Mechanical	9
5.1 Mechanical Outline	9
5.2 PCIe Edge Connector	15
5.3 Platform Design.....	15
6 Thermal	16
6.1 Data Center Environmental Conditions.....	16
6.2 Server Operational Conditions.....	16
6.3 CPU cooling solution.....	17
6.4 M.2 Cooling Solution	18
6.5 Temperature and Power Sensors.....	19
6.6 Server Card Air Baffle:	20
7 Electrical	21
7.1 Design Guidelines.....	21
7.2 Primary X16 Edge Connector A	21
7.3 Extension X16 Edge Connector B	24
7.4 Pin Definitions	27
7.5 Ethernet.....	32
7.6 USB	32
7.7 Serial port	33
7.8 I ² C	33
7.9 NIC Sideband.....	33
7.10 Slot ID and GPIO	33
8 Power	34

8.1	Input	34
8.2	VR Efficiency.....	35
8.3	Input Capacitance.....	35
8.4	Power reading and power capping	35
9	Functional	36
9.1	System on a Chip	36
9.2	Memory	36
9.3	Debug headers	36
9.4	Storage.....	36
9.5	EEPROM	36
9.6	BIOS	37
9.7	Twin Lakes 1S Server Management.....	40
9.8	LEDs	60
10	Environmental Requirements.....	61
10.1	Vibration and Shock	61
11	Prescribed Materials.....	62
11.1	Disallowed Components	62
11.2	Capacitors and Inductors	62
11.3	Component De-rating.....	62
12	Labels and Markings	63
13	Revision History	64

3 Overview

This document describes a single socket server design based on the Intel® Xeon® Processor D-2191, which is referred to hereinafter as Twin Lakes 1S server.

The Twin Lakes 1S server is designed to use Intel® Xeon® Processor D-2191 utilizing the performance and advanced Intelligence of Intel® Xeon® processors packaged into a dense, low-power SoC.

Fueled by Intel® Xeon® Processor D SoC, integrating a multi-core CPU, memory controller, and input/output (IO) device capabilities, the Twin Lakes 1S server is great for single or multi node platforms where a solid balance of performance and IO density are key criteria. However, platform designers must provide adequate power and cooling to properly handle the SoC's power and thermal requirements.

Figure 3-1 illustrates the Twin Lakes 1S server block diagram.

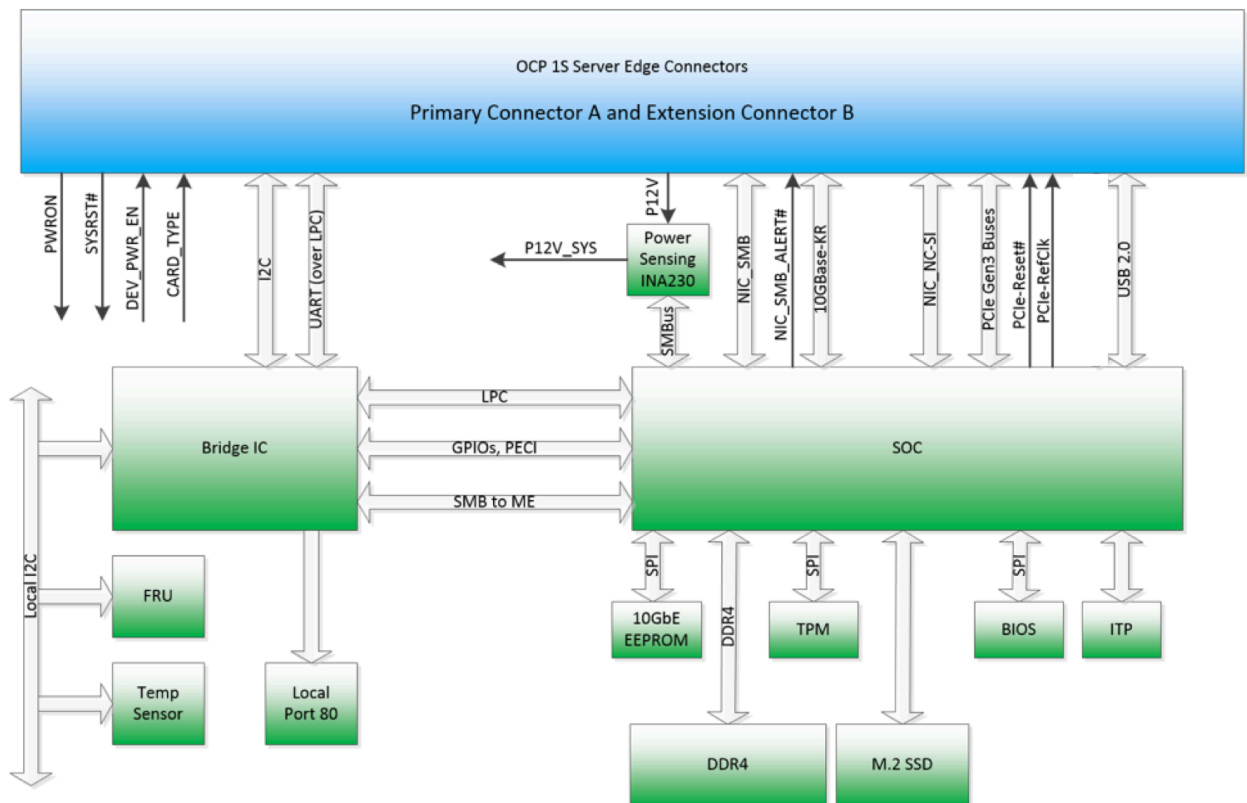


Figure 3-1 Twin Lakes 1S Server Block Diagram

The Twin Lakes 1S server implements primary and extension x16 PCIe edge connectors as defined in the 1S server specification. The primary x16 PCIe edge connector supports:

- PCIe Gen3 ports
- A 10GBase-KR
- A SATA port
- A USB 2.0 port
- A Universal Asynchronous Receiver/Transmitter (UART)

- An I²C bus for server management
- A System Management Bus (SMBus) as the sideband of the integrated network controller

The extension x16 PCIe edge connector supports:

- A Network Controller Sideband Interface (NC-SI) as an alternative sideband of the integrated network controller
- PCIe Gen3 ports

The Twin Lakes 1S server supports three on-card Solid State Drives (SSDs) in the 2280 or 22110 M.2 form factor. One of the M.2 SSD drives is used as boot drive, which is on bottom side of the card next to the SOC. The x4 PCIe link to the boot drive is connected to the PCH portion of the SOC. The boot M.2 drive can be configured to use SATA as the electrical interface through BOM option. Regardless of the electrical interface, the recommended minimum capacity of the boot drive is 256GB. Due to the usage model and capacity of the boot M.2 drive, only 2280 form factor is supported.

The other two M.2 SSD drives are targeted for high performance use case thus only PCIe X4 links from the CPU portion of the SOC are connected to them. Both of them support 2280 or 22110 form factors.

The Twin Lakes 1S server can use either an external network interface controller on the platform through its PCIe interface, or the SoC's own integrated network controller. Generally, the Twin Lakes 1S server assumes a Baseboard Management Controller (BMC) is used on the platform. When the integrated network controller on the SoC is used as a shared NIC, it will support SMBus as the sideband to the BMC on the platform in all power states.

The Twin Lakes 1S server receives 12.5V from the platform with a maximum current load of 7.7A from the primary edge connector and an additional 7.7A from the extension edge connector. The platform, however, defines and controls the maximum power used by the Twin Lakes 1S server. The Twin Lakes 1S server uses the INA230 power sensor at the power input to measure the total card power consumption of the whole server with +/-1% accuracy. The power data measured by this sensor can be used by the platform for power management purposes. When the power consumption has exceeded the power limit, the Bridge IC will generate a throttle event to force the SoC throttle to the lowest power state. The Twin Lakes 1S server supports an Advanced Configuration Power Interface (ACPI)-compliant power button and reset signals from the platform.

A Bridge IC (Texas Instrument's Tiva microcontroller) is used as the management controller on the 1S server and the bridge between the BMC and the SoC. The Bridge IC manages the 1S server on behalf of the BMC on the platform and bridges the BMC and Intel® Xeon® Processor D-2191 SoC's internal management controller Intel® Manageability Engine. To maximize the communication bandwidth between the BMC and the Bridge IC, a dedicated point-to-point I²C bus shall be used.

Twin Lakes 1S server's Field Replaceable Unit (FRU) EEPROM and thermal sensors are connected to the Bridge IC's other I²C buses. There are multiple General Purpose Input/Output connections (GPIOs) between the Bridge IC and Intel® Xeon® Processor D-2191 SoC for error reporting and other management purposes. Intel® Xeon® Processor D-2191 SoC's Intel® Manageability Engine is connected to the Bridge IC via an I²C interface so that the Bridge IC can poll information from the Intel® Manageability Engine. A Low Pin Count (LPC) bus between the Bridge IC and the Intel®

Xeon® Processor D-2191 SoC is connected to enable in-band communications. The BMC can access the Twin Lakes 1S server's thermal sensors, FRU, Intel® Xeon® Processor D-2191 SoC's GPIOs and Intel® Manageability Engine via the Bridge IC with standard IPMI commands.

BIOS, network controller and boot ROM of Bridge IC on the Twin Lakes 1S server, can be updated from in-band connectivity by the SoC or out-of-band thru BMC. While other firmwares (CPLD and VRs) are programmable from out-of-band connectivity by the Bridge IC and BMC.

4 License

© 2016 Facebook.

As of July 26, 2016, the following persons or entities have made this Specification available under the Open Compute Project Hardware License (Permissive) Version 1.0 (OCPHL-P), which is available at <http://www.opencompute.org/.../spec-submission-process/>.

Facebook, Inc.

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 Specification implementer and user assume the entire risk as to implementing or otherwise using the Specification. IN NO EVENT WILL ANY PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS 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 Twin Lakes 1S Server card uses a primary x16 PCIe edge connector and an additional extension x16 PCIe edge connector as the interfaces to the platform.

The overall dimensions of the general card is 210mmx160mm. See Figure 5-1 for the specification drawing including keep-out zones and component restrictions.





Figure 5-1 Twin Lakes 1S Server Mechanical Drawings

The PCB is mounted to a sheet metal carrier which includes flanges on each perpendicular edge that engage the card guides in the chassis. There are no card guide keep-outs on the PCB edge since all available space is needed for traces and components.

The carrier assembly includes 2x ejectors which are used for card injection/ejection into the PCIe connectors. These ejectors rotate horizontally before being allowed to swing open, and they

include finger access cutouts. One ejector claw engages a limit switch, which allows the BMC to detect if the ejector has been opened. The air duct rotates above a fixed support wall.

A high-performance passive cooling solution chosen to cool the SoC effectively.








Figure 5-2 Isometric Views of Twin Lakes 1S Server Assembly








Figure 5-3 Twin Lakes 1S Server Assembly

The thickness of the PCB shall be $1.65 \pm 10\%$ mm to accommodate regular PCIe x16 connectors on platforms. Components of significant height will be placed on the A-side with maximum height limit of 45mm. Low profile components of height less than 2.67mm can be placed on the B-side with exception of M.2s. Boot drive M.2 has the option of low-profile connector or enhance serviceability with 5.8mm connector. High performance M.2 connector shall be 5.8mm connector only.

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 card's edge connector 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.

6 Thermal

6.1 Data Center Environmental Conditions

This section outlines Facebook data center operational conditions.

6.1.1 Location of Data Center/Altitude

Maximum altitude is 6,000 ft above sea level. Any variation of air properties or environmental difference due to the high altitude needs to be deliberated into the thermal design.

6.1.2 Cold-Aisle Temperature

Data centers generally maintain cold aisle temperatures between 18°C and 30°C (65°F to 85°F). The mean temperature in the cold aisle is usually 25°C with 3°C standard deviation. The cold aisle temperature in a data center may fluctuate minutely depending to the outside air temperature. Every component must be cooled and must maintain a temperature below its maximum specification temperature in the cold aisle.

6.1.3 Cold-Aisle Pressurization

Data centers generally maintain cold aisle pressure between 0 inches H₂O and 0.005 inches H₂O. The thermal solution of the system should consider the worst operational pressurization possible, which generally is 0 inches H₂O and 0.005 inches H₂O with a single fan (or rotor) failure.

6.1.4 Relative Humidity

Data centers usually maintains a relative humidity between 20% and 90%.

6.2 Server Operational Conditions

6.2.1 System Volumetric Flow

The unit of airflow (or volumetric flow) used for this spec is cubic feet per minute (CFM). The CFM can be used to determine the thermal expenditure or to calculate the approximate Delta T of the system. The thermal expenditure is quantified by the metric CFM/W, which is calculated by the following formula:

$$\text{Thermal Expenditure} = \frac{\text{System airflow}}{\text{Total system power consumption, including fans}} \quad [\text{CFM/W}]$$

The required airflow is 0.115 airflow per watt in the system level at sea level. The desired airflow per watt is 0.1 or lower up to 45°C (113°F) ambient temperature at sea level.

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 detection of abnormal thermal behaviors in the system. The UCT values for the sensors that are not used in Fan Speed Control

(FSC) should use a 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 a 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) inlet temperature to guarantee the design is free of thermal defect and has high temperature reliability.

6.3 CPU cooling solution

6.3.1 Orientation

The CPU cooling solution must be designed such that performance is not impacted by server card orientation. All vapor chamber, or heat pipe components must operate in any orientation.

6.3.2 Solution Footprint

The CPU cooling solution must fit within the available space on the server card. A figure showing the total available footprint can be seen below in Figure 6-1. Section A in Figure 6-1 is centered over the CPU. The total height of the solution and CPU stack up must be less than 45mm. The CPU cooling solution must not interfere with any tall electrical components on the server card.




Figure 6-1 CPU cooling solution footprint. All dimensions in mm

6.3.3 Passive Cooling Solution

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 6-2. The heat sink may require cutouts in order to avoid tall electrical components on the server card.

6.3.4 Reliability Requirements

The operating life of the server card is a minimum of 3 years with target life of 5 years. All pertinent reliability data for CPU cooling solution must be provided to meet this MTBF.

6.4 M.2 Cooling Solution

Each card must support up to 3 M.2s. Two high power M.2s on the bottom side of the card at the front of the card, and one boot drive M.2 on the bottom side at the rear of the card.

6.4.1 Bottom Side M.2 Cooling Requirements

The cooling solution for M.2 cooling must be thermally optimized design at the lowest cost. Passive cooling is required with a gap pad and heat sink solution being desired. The solution must maintain all M.2 component temperatures within their operational limits during all stress conditions. 5.8mm M.2 connectors must be used to allow spacing for gap pads between M.2 and server card. The total solution stack up must be 12mm or shorter.

6.4.2 Gap Pad Requirements

The gap pads must be soft and elastic as the application will be between two PCBs. The gap pads must not leave any residue when replaced. The Gap pads must have a thermal conductivity of 2 W/m-k or higher. Gap pad thickness is determined by compression vs pressure that gives the lowest loading force while supporting tolerance of components sandwiching the gap pad. Gap pad AVL provided as needed.

6.4.3 Bottom Side boot M.2 Requirements

Bottom side boot M.2 must maintain operation conditions for all components under all stress conditions. The bottom side boot M.2 is downstream from the high speed M.2s in standard orientation and may experience high temperatures due to pre-heating from upstream M.2s. A heat sink with gap pad solution may be required in order to meet bottom side boot M.2 cooling requirements.

6.5 Temperature and Power Sensors

Each card must provide following sensors:

- Temperature sensors for SOC, DIMM, voltage regulators and other critical chips
- Power sensors for the SOC and the whole card
- Voltage sensors for all voltage rails
- One inlet ambient temperature sensor and one outlet ambient temperature sensor

The BMC on the platform must be able to read all these sensors via the Bridge IC. 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 $\pm 2^{\circ}\text{C}$ and desired to be within 2% tolerance across whole operation temperature range. The goal sensor accuracy is $\pm 1^{\circ}\text{C}$.

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




Figure 6-2 Ambient Temperature Sensor Placement Location (Highlighted Areas in Blue)

6.6 Server Card Air Baffle:

The card level air baffle must be designed to help maintain temperatures of all major components on the server card by reducing bypass air, and increasing airflow through key components. The air baffle must be easy to service with the goal of requiring no tooling to remove. The air baffle must not have a large adverse effect on system level pressure drop.

7 Electrical

7.1 Design Guidelines

Refer to Intel® Xeon® Processor D-2191 SoC documents for design guidelines.

7.2 Primary X16 Edge Connector A

The Twin Lakes 1S server uses both primary and extension edge connectors as defined below. There are two changes from 1S server spec 0.5. First, the original KR4 interface on extension edge connector has been reassigned as 7th PCIe Gen3 X4 link, which does not have its own PCIe reference clock and PCIe reset signal. Secondly, 4 of the GPIO pins on primary connector have been redefined with specific purpose now. Last but not least, we also converted the SATA pins to PCIe x1 link to make it possible to add a VGA controller on platform side.

Table 1: Twin Lakes 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	CARD_TYPE
I2C_DATA	6	6	DEV_PWR_EN
GND	7	7	COM_TX
PWR_BTN#	8	8	COM_RX
USB_P	9	9	EJCT_LATCH_DET_N
USB_N	10	10	RESET_BMC
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
PCIE7_TX_P	31	31	GND
PCIE7_TX_N	32	32	GND
GND	33	33	PCIE7_RX_P
GND	34	34	PCIE7_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	KR0_RX_P
GND	46	46	KR0_RX_N
KR0_TX_P	47	47	GND
KR0_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 Twin Lakes 1S server also implements an extension x16 edge connector to bring out additional x16 PCIe lanes and one NC-SI (pin assignments shown in Table 2). This extension X16 Edge connector is referred to as Connector B.

Table 2: Twin Lakes 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
PCIE6_TX0_P	15	15	GND
PCIE6_TX0_N	16	16	GND
GND	17	17	PCIE6_RX0_P
GND	18	18	PCIE6_RX0_N
PCIE6_TX1_P	19	19	GND
PCIE6_TX1_N	20	20	GND
GND	21	21	PCIE6_RX1_P
GND	22	22	PCIE6_RX1_N
PCIE6_TX2_P	23	23	GND
PCIE6_TX2_N	24	24	GND
GND	25	25	PCIE6_RX2_P

GND	26	26	PCIE6_RX2_N
PCIE6_TX3_P	27	27	GND
PCIE6_TX3_N	28	28	GND
GND	29	29	PCIE6_RX3_P
GND	30	30	PCIE6_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 Twin Lakes 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 ² C clock signal. I ² C 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 ² C data signal. I ² C is the primary sideband interface for server management functionality. 3.3VAUX signal. Pull-up is provided on the platform.
I2C_ALERT#	Output	Required	I ² C alert signal. Alerts the BMC that an 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 greater than 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.
CARD_TYPE (GPIO0)	Output	Required	CARD_TYPE is an output signal to inform platform that if this card is a server or a device. For a server, this pin should be tied to GND through a 10K resistor on the 1S server card.
DEV_PWR_EN (GPIO1)	Output	Required	CARD_TYPE is an output signal from a 1S server to enable active power on device side. It is useful when 1S server is going through DC cycling or AC cycling but the server and devices are not in the same power domain. Active high, 3.3VAUX signal.

EJCT_LATCH_DET_N (GPIO2)	Output	Required	EJCT_LATCH_DET_N is an output signal to indicate if the 1S server is fully seated with ejector latch closed and ready for power on. Platform designer can use this signal to control the power to the 1S server and avoid surprise 1S server insertion/removal to/from a hot slot. Active low, 3.3VAUX signal, pull-up should be provided on the platform.
RESET_BMC (GPIO3)	Output	Required	RESET_BMC is an output signal to reset BMC on the platform. Active high, 3.3VAUX signal, pull-down on the platform.
KR0_TX_P/N	Output	Required	Primary 10GBase-KR Ethernet transmit signal. Data is sent from the 1S Server module to the platform.
KR0_RX_P/N	Input	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	PCIe reference clock. These signals may or may not be connected on the platform.
PCIE8_TX_P/N	Output	Required	PCIe Gen2 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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	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	Required	PCIe reference clock. These signals may or may not be connected on the platform.
PCIE6_TX0/1/2/3_P/N	Output	Required	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.
PCIE6_RX0/1/2/3_P/N	Input	Required	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.

PCIE8_RX_P/N	Input	Required	PCIe Gen2 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.
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.

7.5 Ethernet

The Twin Lakes 1S server uses the 10GBase-KR as its optional Ethernet connection. Auto-negotiation and link training are required to work with either the re-timer or the physical layer (PHY) on the platform.

7.5.1 Routing guidelines

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

It is strongly recommended to perform comprehensive signal integrity simulation, analysis and validation for 10GBase-KR 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 USB

The USB connection is USB 2.0.

7.7 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).

7.8 I²C

A single I²C 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 I²C 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.

7.9 NIC Sideband

7.9.1 SMBus

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

7.9.2 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 significantly 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.

7.10 Slot ID and GPIO

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 1S server spec. These pins can be used as GPIO pins if they are not used as slot IDs.

Twin Lakes 1S server redefined all 4 GPIO pins with specific functions. Please refer to pin definition table above for details.

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 Twin Lakes 1S server power capacity

The Twin Lakes 1S server's maximum power is 180W with full configuration. It is critical to develop a sophisticated thermal solution for the platform to keep the Twin Lakes 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 Twin Lakes 1S server card from the platform perspective. However, a standby 3.3V_AUX power rail on the card is required to power the Bridge IC at all power states. The SoC also requires some stand-by power rails. For example, the SoC's integrated network controller requires stand-by power rails to keep the out-of-band function alive when the SoC is in standby mode. A power CPLD that controls the power sequencing of the entire server is also powered by the 3.3V_AUX standby power rail.

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 leakage path among the power domains including 1S server stand-by, 1S server core, and baseboard platform power.

8.1.3 Standby power budget

The Twin Lakes 1S server consumes minimal power when it operates in standby mode.

8.2 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.

8.3 Input Capacitance

The capacitance on the input 12V rail of the Twin Lakes 1S server shall be optimized. Platform designers must ensure overall capacitance on the 12.5V rail of the entire platform meets the system's power supply requirement and does not cause instability.

8.4 Power reading and power capping

The Twin Lakes 1S server implements sophisticated power management features.

The Twin Lakes 1S server shall have the power monitoring capability to read power consumption reliably and accurately and can report a one-second average power reading with 3% accuracy. As shown in figure 3-1, a power sensor with I²C interface is used at the 12.5V power input to the card and is readable by the Bridge IC. Both the BMC and SoC can obtain the whole-card power consumption information from the Bridge IC.

The Twin Lakes 1S server can throttle itself down to lowest possible power state as quickly as possible when the platform asserts the FAST_THROTTLE_N signal or it receives request from BMC, or over power event reported by on-card power monitor.

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 Twin Lakes 1S server uses an Intel® Xeon® Processor D-2191 SoC that incorporates a multi-core CPU, memory controller, and other IO devices.

9.2 Memory

The Twin Lakes 1S server uses DDR4 memory bus.

9.3 Debug headers

The Twin Lakes 1S server shall support traditional ITP/XDP debug headers.

9.4 Storage

The Twin Lakes 1S server supports three M.2 solid-state drives in 2280 or 22110 form factors. Boot M.2 slot is only available in 2280 form factor and it can be configured as either SATA or PCIe interface through BOM options, but not both. A minimum 256GB M.2 SATA or NVMe SSD is required as a boot device and for logging purpose. Two additional SSD drives are designed in to support applications that require high disk performance. These two M.2 slots only support PCIe X4 links but both support SSD drives in 2280 or 22110 form factors.

9.5 EEPROM

The Twin Lakes 1S server includes a 128Kbits I²C-accessible Electrically Erasable Programmable Read-Only Memory (EEPROM). The EEPROM is accessible from the platform via the Bridge IC. The EEPROM contains 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 entries:

- 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 T_{jMAX} (Maximum Junction Temperature)

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 (like 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 and 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 is logged by the BIOS or 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 others.
- 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.

9.7 Twin Lakes 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 I²C bus as the management interface. This section identifies the required information that must be accessible from the BMC.

9.7.1 Bridge IC

The Twin Lakes 1S server uses a Bridge IC is defined as the bridging device between the SoC and the BMC. The Bridge IC is 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 I²C bus. To enable prompt communication, this I²C 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 has FRU EEPROM and thermal sensors on a local I²C bus. The BMC can communicate with the Bridge IC to inquire the FRU and thermal data through the Intelligent Platform Management Bus (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.

The Bridge IC has a dedicated I²C bus to the SoC's Intel® Manageability Engine that supports IPMB. The Bridge IC now behaves as a transparent bridge to forward IPMI messages between the BMC and SoC's Intel® Manageability Engine. With this transparent bridge, the BMC can directly work with SoC's Intel® Manageability Engine to perform most of server management functions.

The Bridge IC also implements a system interface LPC to enable in-band server manageability. A KCS interface is implemented and it supports both standard SMM and SMS interfaces.

The Bridge IC monitors the Twin Lakes 1S server's sensors, such as voltage sensors, power sensors, and digital sensors for critical GPIOs. The BMC can 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 must have two Serial Peripheral Interfaces (SPI). The first SPI bus is used as its own boot ROM if it does not have an integrated boot ROM. The second SPI can be used to reprogram the SoC's boot ROM when it is corrupted with a multiplexer and BMC's support. The Bridge IC on the Twin Lakes 1S server is a Texas Instrument's Tiva microcontroller.

9.7.2 I²C addressing

The Twin Lakes 1S server and BMC communicates using IPMI 2.0 commands transmitted over the I²C connection through a Bridge IC on the 1S Server card. The I²C bus address for the Bridge IC is 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 provides ways to transfer messages between them via KCS interfaces.

For in-band management, the Bridge IC can 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 I²C bus meant for the SoC's Intel® Manageability Engine, the Bridge IC shall forward the command on the I²C bus and send the 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 I²C. When the BMC sends SOL data, it shall be emitted via the serial port.

9.7.4 Platform Discovery and Configuration

The Bridge IC provides a way for the BMC to discover platform capabilities such as electrical interface assignment. It also provides 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 the 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 the Bridge with the latest POST Code.

9.7.6 IPMB Interface

The Bridge IC provides an IPMB interface for the BMC to access various IPMI resources on the Twin Lakes 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 can update firmware of the programmable devices on the Twin Lakes 1S Server, such as BIOS, the SoC's Intel® Manageability Engine firmware, the CPLD image, various VR firmware and the Bridge IC's firmware.

Bridge IC provides a way for the BMC to access the version information and initiate update process of various firmware components on the Twin Lakes 1S Server, as well as 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 but not on the Twin Lakes 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 I²C bus between them.

The BMC can then use the LED information to control the 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 a GPIO register block. In this way, the BMC can control the GPIO behind the Bridge IC (or this hardware abstraction layer) through accessing this register block.

The Bridge IC shall provide a way for BMC to configure GPIO pin direction, interrupt capability, and provide a way to get/set the current status of GPIO signals. It shall send an interrupt message to the BMC when the interrupt enabled GPIO signal changes its 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 4.

Table 4: Bridge IC GPIO Table

No.	GPIO offset	GPIO Pin Function	Comments
1	Byte 1 - bit [0]	Power Good – CPU Core	Indicates that the CPU's core power input is good
2	Byte 1 - bit [1]	Power Good – PCH core	Indicates that the PCH's core power input is good
3	Byte 1 - bit [2]	DDR Channel A/B Voltage regulator hot	Indicates that the DDR Channel A/B Voltage Regulator is hot
4	Byte 1 - bit [3]	DDR Channel D/E Voltage regulator hot	Indicates that the DDR Channel D/E Voltage Regulator is hot
5	Byte 1 - bit [4]	CPU VccIN Voltage regulator hot	Indicates that the CPU VccIN Voltage Regulator is hot
6	Byte 1 - bit [5]	CPU Throttle	System firmware request CPU throttle
7	Byte 1 - bit [6]	PCH Hot	Indicates PCH temperature is over setpoint
8	Byte 1 - bit [7]	DIMM Hot	Indicates that one or more DIMM is hot
9	Byte 2 - bit [0]	CPU thermal trip	Indicates that CPU experienced over temperature event and shut-down
10	Byte 2 - bit [1]	PCH thermal trip	Indicates that PCH experienced over temperature event and shut-down
11	Byte 2 - bit [2]	CPU FIVR fault	Indicates a CPU internal Voltage Regulator Error condition
12	Byte 2 - bit [3]	CPU Catastrophic Error	Indicates that the CPU experienced catastrophic error
13	Byte 2 - bit [4]	CPU Non-Recoverable Error	Indicates that the CPU experienced non-recoverable error
14	Byte 2 - bit [5]	CPU Critical Error	Indicates that the CPU experienced critical error

15	Byte 2 - bit [6]	CPU Non-Critical Error	Indicates that the CPU experienced non-critical error
16	Byte 2 - bit [7]	Sleep S4 state	When low, indicates CPU has entered S4 state or lower
17	Byte 3 - bit [0]	Non-maskable interrupt	Non-maskable interrupt
18	Byte 3 - bit [1]	System management interrupt	System management interrupt
19	Byte 3 - bit [2]	Platform Reset	Platform and PCIe Reset
20	Byte 3 - bit [3]	Front Panel Reset Input	Initiate platform reset (front panel reset button being pressed)
21	Byte 3 - bit [4]	Front Panel Reset Output	Host Reset Output from Bridge-IC
22	Byte 3 - bit [5]	Bios Power-on Self Test (POST) complete	Bios Power-on Self Test (POST) complete
23	Byte 3 - bit [6]	Sleep S3 state	When low, indicates CPU has entered S3 state or lower
24	Byte 3 - bit [7]	Power Good – CPU VccIN	Indicates that the VCCIN Voltage Regulator power is good
25	Byte 4 - bit [0]	Boot SPI selection	Select boot SPI0 or SPI1 as boot SPI
26	Byte 4 - bit [1]	Ejector latch detection	Identify ejector fully closed
27	Byte 4 - bit [2]	BMC Reset	Reset BMC from host
28	Byte 4 - bit [3]	At-scale-debug (ASD) TCK selection	Select at-scale-debug (ASD) TCK drive CPU TCK or PCH TCK
29	Byte 4 - bit [4]	BMC ready	Indicated BMC is ready
30	Byte 4 - bit [5]	Host / Bridge-IC UART select	Select Host UART or Bridge-IC Debug UART
31	Byte 4 - bit [6]	I2C MUX reset	Reset I2C MUX
32	Byte 4 - bit [7]	At-scale-debug (ASD) PREQ	CPU probe mode request
33	Byte 5 - bit [0]	At-scale-debug (ASD) JTAG TRST	JTAG Reset
34	Byte 5 - bit [1]	System Throttle	Indicates system is currently throttling
35	Byte 5 - bit [2]	At-scale-debug (ASD) PRDY	CPU probe mode ready
36	Byte 5 - bit [3]	XDP Present	Indicates traditional ITP connected
37	Byte 5 - bit [4]	ASD Present	Indicates at-scale-debug connected
38	Byte 5 - bit [5]	CPU power debug	Use for debug CPU integrated VR
39	Byte 5 - bit [6]	JTAG MUX selection	Select JTAG connection to ITP or ASD

9.7.10 IPMI commands

The Bridge IC must support the IPMI commands shown in Table 5.

Table 5: 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 ² C	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
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 mode	OEM (0x38)	13h
Set VR monitor Enable	OEM (0x38)	14h
Get VR monitor Enable	OEM (0x38)	15h
Reset BMC	OEM (0x38)	16h
Read BIOS image	OEM (0x38)	18h
Get Flash size	OEM (0x38)	19h
Set Pump Duty	OEM (0x38)	1Bh
Get Pump Duty	OEM (0x38)	1Ch
Set BIOS Chip Select	OEM (0x38)	1D
Get BIOS Chip Select	OEM (0x38)	1F
Set JTAG Tap State	OEM (0x38)	21
Shift JTAG Data	OEM (0x38)	22
Set System GUID	OEM (0x38)	EFh

Table 6 provides details of the IPMI Original Equipment Manufacturer (OEM) commands that are defined Table 6.

Table 6: Bridge IC supported IPMI command Table

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
01h	Send request message to BMC	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – Request interface</p> <p>01h: Intel® Manageability Engine</p> <p>02h: SOL</p> <p>03h: KCS SMS</p> <p>04h: KCS SMM</p> <p>Byte 5:X – Request data</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5 – Request interface</p> <p>01h: Intel® Manageability Engine</p> <p>02h: SOL</p> <p>03h: KCS</p> <p>Byte 6:X – Response data</p>	<p>This command is used for Bridge IC transfer request to BMC.</p> <p>For example:</p> <ol style="list-style-type: none"> 1. Bridge IC gets “Get Device ID” command 0x06 0x01 from KCS. 2. Bridge IC will send this command to BMC as below. 0x38 0x01 0x03 0x06 0x01 3. BMC responds Get Device ID data.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
02h	Send request message to Bridge IC	<p>Request: Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 –Receive interface</p> <p>01h: Intel® Manageability Engine</p> <p>02h: SOL</p> <p>Byte 5:X – Request data from BMC</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5 –Receive interface</p> <p>01h: Intel® Manageability Engine</p> <p>02h: SOL</p> <p>Byte 6:X – Response data</p>	<p>This command is used for BMC send request to Bridge IC.</p> <p>For example:</p> <ol style="list-style-type: none"> 1. When BMC want to send “Get Device ID” command to ME. It can use this command: 0x38 0x02 0x01 0x06 0x01 2. When Bridge IC receive this command, It will send “Get Device ID” command to Intel® Manageability Engine and get the response from Intel® Manageability Engine . 3. Bridge IC responds this command to BMC.
03h	Get all GPIO status	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5:8 – Get all GPIO status</p> <p>0b: Low</p> <p>1b: High</p>	<p>This command used by BMC to get GPIO status from Bridge IC. Refer to Table 8 GPIO mapping table.</p>

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
04h	Set all GPIO status	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4:7 – GPIO enable mask, refer to GPIO mapping table</p> <p>0b: Disable</p> <p>1b: Enable</p> <p>Byte 8:11 – Set all GPIO status</p> <p>0b: Low</p> <p>1b: High</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command used by BMC to set GPIO status from Bridge IC. Refer to Table 8 GPIO Mapping Table.
05h	Get GPIO configuration	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4:7 – GPIO enable mask, refer to GPIO mapping table</p> <p>0b: Disable</p> <p>1b: Enable</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5:X – GPIO configuration (one byte for one GPIO pin configuration)</p> <p>Bit[0] – Input/output pin</p> <p>0b : Input pin</p> <p>1b : Output pin</p> <p>Bit[1] – interrupt disable/enable</p> <p>0b : Disable</p> <p>1b : Enable</p> <p>Bit[2] – Edge trigger</p> <p>0b : Edge trigger (default)</p> <p>Bit[3:4] – Trigger type</p> <p>00b : Falling edge</p> <p>01b : Rising edge</p> <p>10b : Both</p> <p>11b : Reserved</p>	This command used by BMC to get GPIO configuration from 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	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4:7 – GPIO enable mask, refer to GPIO mapping table</p> <p>0b: Disable</p> <p>1b: Enable</p> <p>Byte 8:X – GPIO configuration (one byte for one GPIO pin configuration)</p> <p>Bit[0] – Input/output pin</p> <p>0b : Input pin</p> <p>1b : Output pin</p> <p>Bit[1] – interrupt disable/enable</p> <p>0b : Disable</p> <p>1b : Enable</p> <p>Bit[2] – Edge trigger</p> <p>0b : Edge trigger (default)</p> <p>Bit[3:4] – Trigger type</p> <p>00b : Falling edge</p> <p>01b : Rising edge</p> <p>10b : Both</p> <p>11b : Reserved</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command used by BMC to set GPIO configuration to Bridge IC. Refer to Table 8 GPIO Mapping Table.
07h	Send interrupt to BMC	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – Interrupt GPIO number, refer to GPIO mapping table</p> <p>Byte 5 – Trigger type</p> <p>00h: Falling edge</p> <p>01h: Rising edge</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command used for Interrupt notification from Bridge IC sends to 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 – Port 80 data Response: Byte 1 – Completion Code 00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first	Bridge IC support maximum 230 bytes to buffer BIOS POST Code when BMC is not ready. The POST Code data will be in FIFO manner i.e. with the first POST Code as the first byte. But in case BMC is ready, 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 – Port 80 data Byte 2:4 – IANA ID – 00A015h, LS byte first	BMC can get all POST Code data by this command. 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. Bridge IC clear buffer when system power on. The maximum buffer data length is 230 bytes.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
09h	Firmware Update	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – update target</p> <p>00h: BIOS</p> <p>01h: CPLD</p> <p>- bit[7] = 1, last package for image data</p> <p>02h: Bridge IC boot loader from OOB</p> <p>- bit[7] = 1, last package for image data</p> <p>03h: Bridge IC boot loader from In-Band</p> <p>- bit[7] = 1, last package for image data</p> <p>04h: VR</p> <p>- bit[7] = 1, last package for image data</p> <p>Byte 5:8 – Offset</p> <p>Byte 9:10 – Data length</p> <p>Byte 11:X – Update image data</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>80h – Write flash error</p> <p>81h – Power status check fail</p> <p>82h – Data length error</p> <p>83h – Flash erase error</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command is used to update BIOS and CPLD Firmware from BMC.
0Ah	Firmware Verify	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – update target</p> <p>00h: BIOS</p> <p>01h: CPLD</p> <p>02h: Bridge IC boot loader</p> <p>03h: VR</p> <p>Byte 5:8 – Offset</p> <p>Byte 9:12 – Data length</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>80h – Checksum error</p> <p>82h – Data length error</p> <p>84h – Read flash error</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5:8 – Checksum</p>	This command is used to verify BIOS and CPLD Firmware from BMC.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
0Bh	Get Firmware version	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – update target</p> <p>01h: CPLD</p> <p>02h: Bridge IC</p> <p>03h: Intel® Manageability Engine version</p> <p>04h: Bridge IC Bootloader</p> <p>05h: VCCIO VR</p> <p>06h: VCCIN VR</p> <p>07h: VCCSA VR</p> <p>08h: DDR_AB VR</p> <p>09h: DDR_DE VR</p> <p>0Ah: VNNPCH VR</p> <p>0Bh: 1V05PCH VR</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5:X –</p> <p>CPLD version (Hexadecimal) – CPLD user code 4 bytes.</p> <p>Bridge IC version (Decimal) – 2 bytes length, ex: 1.03. Return data will be 0x01 0x03.</p> <p>Intel® Manageability Engine version (Decimal) – 5 bytes length, ex: Intel® version 03.0.0.010. Return data will be 0x03 0x00 0x00 0x01 0x00.</p> <p>VR version (Hexadecimal) – VR user data, 4 bytes length.</p> <p>Ex: 3d 01 11 00, user data 0 : 0x3d01, user data 1 : 0x1100</p> <p>Bridge IC bootloader version (Decimal) – 2 bytes length, version 1.08 Return data will be 0x01 0x08</p>	
0Ch	Enable Bridge IC update flag	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – Enable update interface flag</p> <p>00h: UART</p> <p>01h: I2C</p> <p>02h: LPC</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command is used to enable Bridge IC update flag from BMC.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
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 2:4 – IANA ID – 00A015h, LS byte first Byte 5 – LED frequency 00h: No blinking 01h: Solid on 02h: Slow flashing 03h: Fast flashing	
0Eh	Get Bridge IC configuration	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 2:4 – IANA ID – 00A015h, LS byte first Byte 5 – Bit[0] – SOL interface 0b : Disable 1b : Enable Bit[1] – Port 80 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	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p> <p>Byte 5 – 0x03: A13/A14 PCIe0 RefClk</p> <p>Byte 6 – 0x02: A17/A18 PCIe0 Lane 0, Gen3</p> <p>Byte 7 – 0x02: A21/A22 PCIe0 Lane 1, Gen3</p> <p>Byte 8 – 0x02: A25/A26 PCIe0 Lane 2, Gen3</p> <p>Byte 9 – 0x02: A29/A30 PCIe0 Lane 3, Gen3</p> <p>Byte 10 – 0x05: A33/A34 SATA0, Gen3</p> <p>Byte 11 – 0x03: A37/A38 PCIe2 RefClk</p> <p>Byte 12 – 0x02: A49/A50 PCIe1 Lane 0 Gen3</p> <p>Byte 13 – 0x02: A53/A54 PCIe1 Lane 1 Gen3</p> <p>Byte 14 – 0x02: A57/A58 PCIe1 Lane 2 Gen3</p> <p>Byte 15 – 0x02: A61/A62 PCIe1 Lane 3 Gen3</p> <p>Byte 16 – 0x02: A65/A66 PCIe2 Lane 0 Gen3</p> <p>Byte 17 – 0x02: A69/A70 PCIe2 Lane 1 Gen3</p> <p>Byte 18 – 0x02: A73/A74 PCIe2 Lane 2 Gen3</p> <p>Byte 19 – 0x02: A77/A78 PCIe2 Lane 3 Gen3</p> <p>Byte 20 – 0x02: B15/B16 PCIe0 Lane 0, Gen3</p> <p>Byte 21 – 0x02: B19/B20 PCIe0 Lane 1, Gen3</p> <p>Byte 22 – 0x02: B23/B24 PCIe0 Lane 2, Gen3</p> <p>Byte 23 – 0x02: B27/B28 PCIe0 Lane 3, Gen3</p> <p>Byte 24 – 0x05: B31/B32 SATA0, Gen3</p> <p>Byte 25 – 0x03: B35/B36 PCIe1 Ref Clk</p> <p>Byte 26 – 0x02: B51/B52 PCIe1 Lane 0 Gen3</p> <p>Byte 27 – 0x02: B55/B56 PCIe1 Lane 1 Gen3</p> <p>Byte 28 – 0x02: B59/B60 PCIe1 Lane 2 Gen3</p> <p>Byte 29 – 0x02: B63/B64 PCIe1 Lane 3 Gen3</p> <p>Byte 30 – 0x02: B67/B68 PCIe2 Lane 0 Gen3</p> <p>Byte 31 – 0x02: B71/B72 PCIe2 Lane 1 Gen3</p> <p>Byte 32 – 0x02: B75/B76 PCIe2 Lane 2 Gen3</p> <p>Byte 33 – 0x02: B79/B80 PCIe2 Lane 3 Gen3</p>	<p>For scalability needs, we list all configurations in this command response. If we store that information in BMC FW and check by FRU ID, we will need to modify BMC FW whenever there is a new card. We recommend storing the info on MB.</p>

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
10h	Set Bridge IC configuration	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – Config Bridge IC</p> <p>Bit[0] – SOL interface</p> <p>0b : Disable</p> <p>1b : Enable</p> <p>Bit[1] – Port 80</p> <p>0b : Disable, Bridge IC will not send post code to BMC</p> <p>1b : Enable</p> <p>Bit[2] – KCS</p> <p>0b : Disable, Bridge IC will not send KCS command to BMC</p> <p>1b : Enable</p> <p>Bit[3] – IPMB message</p> <p>0b : Disable, Bridge IC will not send IPMB message to BMC</p> <p>1b : Enable</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	When BMC enter update mode, BMC can use this command to disable various communication to Bridge IC.
11h	Bridge IC reset cause	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 4 – Reset cause</p> <p>0x00 : Cold reset by Firmware update</p> <p>0x01 : Watchdog timeout</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	Bridge IC will send this command to notify BMC when Bridge IC is reset.
13h	Bridge IC enter update mode	<p>Request:</p> <p>Byte 1:3 – IANA ID – 00A015h, LS byte first</p> <p>Byte 2 – Firmware mode</p> <p>0x01 : normal mode</p> <p>0x0F : update mode</p> <p>Response:</p> <p>Byte 1 – Completion Code</p> <p>00h - Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2)</p> <p>Byte 2:4 – IANA ID – 00A015h, LS byte first</p>	This command is used to notify BMC, Bridge IC enter update mode and normal mode before and after Firmware update.

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
14h	Set VR Monitor Enable	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Enable/Disable VR Monitor 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) Byte 2:4 – IANA ID – 00A015h, LS byte first	
15h	Get VR Monitor Enable	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 2:4 – IANA ID – 00A015h, LS byte first Byte 5 – VR Monitor 0x00: Disable 0x01: Enable	
16h	Reset BMC	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 2:4 – IANA ID – 00A015h, LS byte first	
18h	Read BIOS image	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – update target 00h: BIOS Byte 5:8 – Offset Byte 9 – Data length Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) 81h – Power status check fail Byte 2:4 – IANA ID – 00A015h, LS byte first Byte 5:N – BIOS image data	

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
19h	Get Flash Size	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – target 00h: BIOS Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first Byte 5:8 – flash size	
1Bh	Set Pump Duty	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Duty Value, range from 0 to 100 Byte 5 – Pump number, in this project it's set to 0 Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first	
1Ch	Get Pump Duty	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Pump number, in this project it's set to 0 Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first Byte 5 – Pump Duty	
1Dh	Set BIOS Chip Select	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – BIOS chip select index 0x00: Select SPI 0 0x01: Select SPI 1 Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first	

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
1Fh	Get BIOS Chip Select	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – Pump number, in this project it's set to 0 Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first Byte 5 – BIOS Chip Select Index 0x00: select SPI 0 0x01: select SPI 1	
21h	Set JTAG Tap State	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4 – JTAG TMS Bit length Byte 5 – JTAG TMS Bit data Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first	
22h	Shift JTAG Data	Request: Byte 1:3 – IANA ID – 00A015h, LS byte first Byte 4:5 – write data bit length Byte 5:n – write data Byte n:n+1 – read data bit length Byte n+2 – last transaction index 0x00 – not the last transaction, keep the TMS state 0x01 – last transaction, set tap state to Exit1 state Response: Byte 1 – Completion Code 00h – Success (Remaining standard Completion Codes are shown in IPMI spec v2.0 table 5-2) Byte 2:4 – IANA ID – 00A015h, LS byte first Byte 5:n – JTAG TDO response data	

Net Function = OEM (0x38), LUN = 00			
Code	Command	Request, Response Data	Description
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 serial console of the Twin Lakes server is 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 BMC must be configurable to enable either an immediate power-on, 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 T_{jMAX}
- 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

Along the top edge of the card, a blue LED is used to indicate 12V status of the server.

There is also a blinking amber heartbeat LED on the Twin Lakes 1S server to indicate that the Bridge IC is in operating mode.

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: -40°C to +70°C (short-term storage) *
- Operating altitude with no de-rating to 6000 ft

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 7: Vibration and Shock Requirements

	Operating	Non-Operating
Vibration	0.3G, 5 to 500 to 5 Hz per sweep, 10 sweeps at 1 octave/minute, test along three axes 5-20Hz – 6db/Oct 20-200Hz – 0.0003 G ² /Hz 200-500 – -6db/Oct	1G, 5 to 500 to 5 Hz per sweep, 10 sweeps at 1 octave/minute, test along three axes
Shock	6G, half sine, 11ms, 5 shocks, test along three axes	12G, half sine, 11ms, 10 shocks, test along 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 2 Directive (2011/65/EU)**
- 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 X5R or better (C0G 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, derating analysis is based on at least 20% derating.

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	Silkscreen	No
UL Marking	Silkscreen	No

13 Revision History

Author	Description	Revision	Date
Yan Zhao	▪ Initial draft.	0.1	9/12/2016
Jon Ehlen			
Jarrold Clow			
Sai Dasari			
Yan Zhao	▪ Updated with Intel's comments.	0.2	1/28/2017
Jon Ehlen	▪ Updated with latest mechanical details.	0.3	1/30/2017
Yan Zhao			
	▪ Updated License to OCPHL-P 1.0.		
Jarrold Clow	▪ Updated thermal and environmental requirements sections.	0.4	2/1/2017
Damien Chong	▪ Update with latest requirements	0.9	11/2/2017
Damien Chong	▪ Update with latest S&V requirements	0.95	11/20/2017
	▪ Update document formatting		
Damien Chong	▪ Fix typo	0.96	12/22/2017
	▪ Update RoHS requirements		
Damien Chong	▪ Update document per legal & comms feedback	0.97	1/4/2018
Damien Chong	▪ Update SoC name to Xeon D-2191	1.00	2/26/2018
	▪ Fix formatting issues		

