

OPEN

Compute Project

Open edge server

Revision 0.1

Author: Samuli Toivola, Nokia

1 License (OCP CLA Option)

Contributions to this Specification are made under the terms and conditions set forth in **Open Compute Project Contribution License Agreement (“OCP CLA”) (“Contribution License”)** by:
Nokia Solutions and Networks OY

You can review the signed copies of the applicable Contributor License(s) for this Specification on the OCP website at <http://www.opencompute.org/products/specsanddesign>

Usage of this Specification is governed by the terms and conditions set forth in **Open Compute Project Hardware License – Permissive (“OCPHL Permissive”), (“Specification License”)**.

You can review the applicable Specification License(s) executed by the above referenced contributors to this Specification on the OCP website at <http://www.opencompute.org/participate/legal-documents/>

Note: The following clarifications, which distinguish technology licensed in the Contribution License and/or Specification License from those technologies merely referenced (but not licensed), were accepted by the Incubation Committee of the OCP: **[insert “None” or a description of the applicable clarifications].**

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Contents

1	License (OCP CLA Option)	2
2	Introduction	4
3	Production Architecture Overview	4
3.1	Server Board Block Diagram	7
3.2	Server Board	8
3.2.1	Server Board Placement	8
3.2.2	Server Board Dimension	9
3.2.3	Processor	10
3.2.4	PCH	10
3.2.5	Memory	11
3.2.6	PCIE Lanes	12
3.2.7	SATA	13
3.2.8	USB	14
3.2.9	AST2500	15
3.2.10	TPM	16
3.2.11	UART	16
3.2.12	LOM	16
3.2.13	Power connector for Accelerator Card	17
3.2.14	Thermal Design	17
4	Chassis Spec	20
4.1	3U System Overview	20
4.1.1	3U Chassis Dimension	21
4.1.2	2U Sled Overview	21
4.1.3	1U Sled Overview	24

2 Introduction

Nokia has designed New server concept to fulfill various Edge computing requirements and use cases.

Open Edge Server will be single-socket server barebone that support the Purley Platform Xeon-SP processors in combination with the Lewisburg PCH (PCH) to provide a balanced feature set between technology leadership and cost.

This specification defines the interfaces and connection topology of various buses and control signals in the compute node.

3 Production Architecture Overview

Open Edge Server main characteristics and features are listed below:

Table 1: Server Feature List

Board Name	Open Edge Server Mother Board
Form Factor	Height: 1U tall Width: 215mm Compute Node support: 1/2 width sleds Node count: up to five 1U sleds or two 2U sleds in 3U sub chassis
Mother Board size	W 205.8mm x L 407.95mm, 10 layers, 1.8mm, 8 DIMMs
CPU	Intel Purley platform/ Skylake or Cascade Lake-SP;
Max Processor Wattage	250W
Chipset	Intel Lewisburg PCH C621 or C627
Memory	DIMM Socket Count: 8 Total, CH A/B/D/E 1DPC; CH C/F 2DPC (Support AEP) DIMM Types: DDR4 RDIMM, 1.2V, 2133/2400/2666/2933 DIMM Capacities: 16GB, 32GB,64GB UP to 384GB (64Gx6) of memory for LRDIMM UP to 384GB (64Gx6) of memory for RDIMM

<p>PCIe Expansion Slot</p>	<p>Type: CPU native PCIe Gen3 Port PE1 ABCD: x16 PCIe link to OCP CONN A/B Port PE3 ABCD: x16 PCIe link to 200pin riser CONN Port PE2 AB: x8 PCIe link to 160pin riser CONN Port PE2 CD: x8 PCIe link to 200pin riser CONN DMI: PCIe link x4 to PCH</p> <p>Type: Lewisburg PCH DMI: PCIe link x4 to CPU PCIe link x1 to I210 PCIe link x1 to BMC PCIe link x2 to M.2 riser CONN PCIe link x8 to 160pin riser CONN</p>
<p>Front Side IO</p>	<p>(1) USB 3.0 port (1) Mini USB port for debugging (1) Power Button with LED include identification function (1) Reset Button (1) System status LED (1) HDD Activity LED (1) BMC Heartbeat LED</p>
<p>Network</p>	<p>LOM: Intel® Ethernet Controller I210 (I210) Lewisburg KRx4 Integrated Network Solution with OCP PHY board (Optional)</p>
<p>Video</p>	<p>ASPEED AST2500 8MB DDR4 video memory</p>
<p>FAN</p>	<p>1U Sled Fans: 4028 fan Count: 4</p> <p>2U Sled Fans: 8056 fan Count: 2</p>
<p>ACPI</p>	<p>ACPI compliance, S0, S5 support. (* No S1 and S3 support.)</p>
<p>Power-Supply</p>	<p>2000 Watt DC to DC power supply/2000 Watt AC to DC power supply</p>
<p>TPM</p>	<p>TPM 2.0 (Nuvoton)</p>
<p>1U PCIe Expansion Slot</p>	<p>1U Sled – (1) OCP 2.0 Mezz x16 (Switch to PCH if not using OCP mezz)</p>

	<ul style="list-style-type: none"> (1) Riser card connectivity <ul style="list-style-type: none"> (1) PCIe x16 normal card-edge connector (2) U.2/SATA disk board connectors
<p>2U PCIe Expansion Slot</p>	<p>2U Sled – Option#1</p> <ul style="list-style-type: none"> (1) OCP 2.0 Mezz x16 (Switch to PCH if not using OCP mezz) (1) Riser card connectivity <ul style="list-style-type: none"> (1) PCIe x16 normal card-edge connector (2) PCIe x8 normal card-edge connectors (4) U.2/SATA disk board connectors (1) M.2 Riser connector <p>2U Sled – Option#2</p> <ul style="list-style-type: none"> (1) OCP 2.0 Mezz x16 (Switch to PCH if not using OCP mezz) (1) Riser card connectivity <ul style="list-style-type: none"> (1) PCIe x16 normal card-edge connector (1) PCIe x8 normal card-edge connector (4) U.2/SATA disk board connector (1) M.2 Riser connector
<p>Chassis</p>	<p>Operating temperature: -5°C to 45°C</p> <p>Short term operating temperature: -5°C to 55°C</p> <p>Non-operating temperature -40°C to 70°C</p> <p>Operating relative humidity 5% to 95% RH</p> <p>Non-operating relative humidity 5% to 95%RH</p>

3.1 Server Board Block Diagram

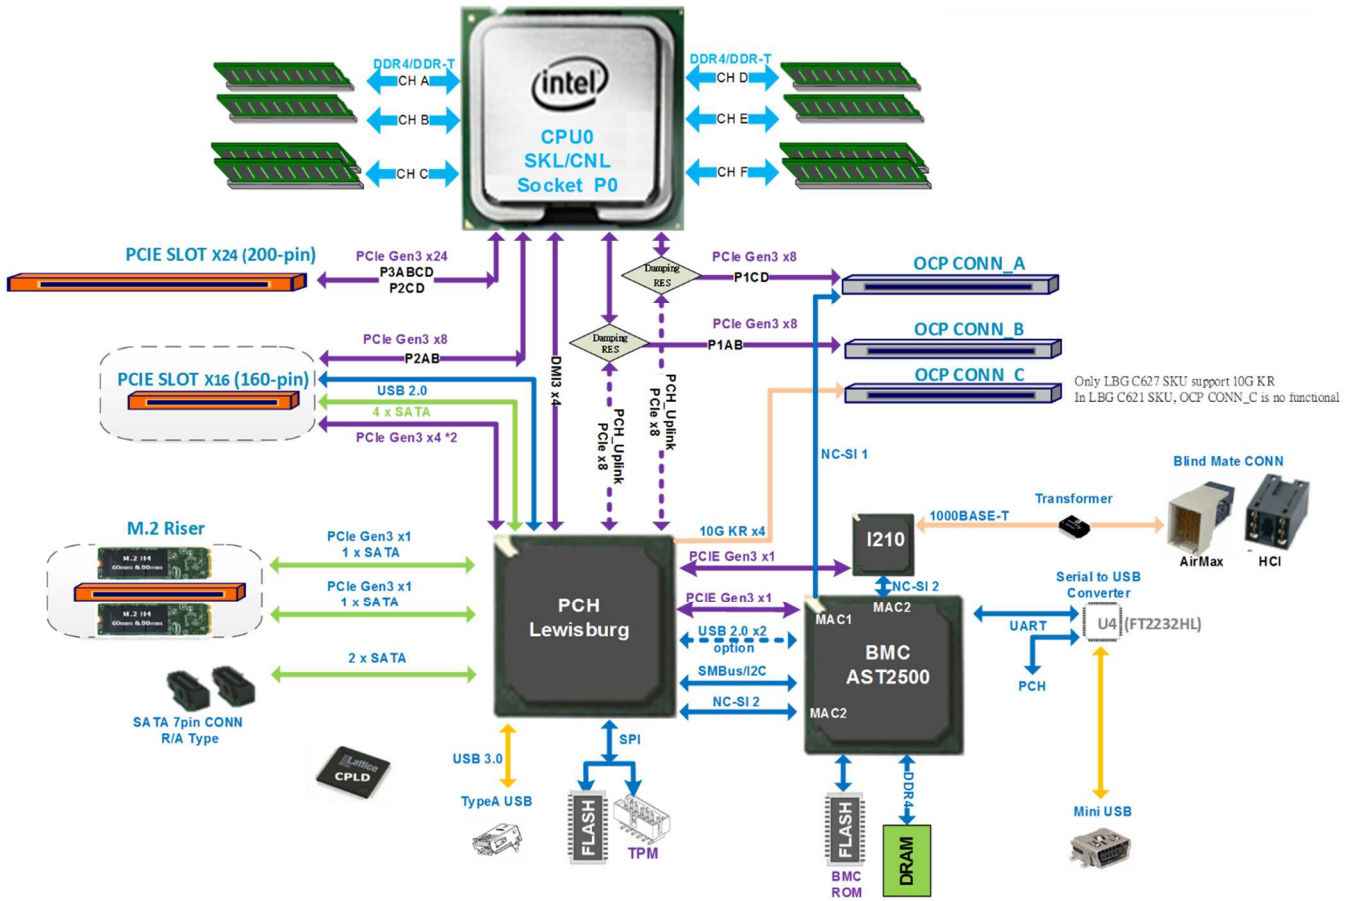


Figure 1: Server Board Block Diagram

3.2 Server Board

3.2.1 Server Board Placement




Figure 2: Server Board Placement

3.2.2 Server Board Dimension

Server Board dimension is 205.8mm x 407.95mm (WxL)

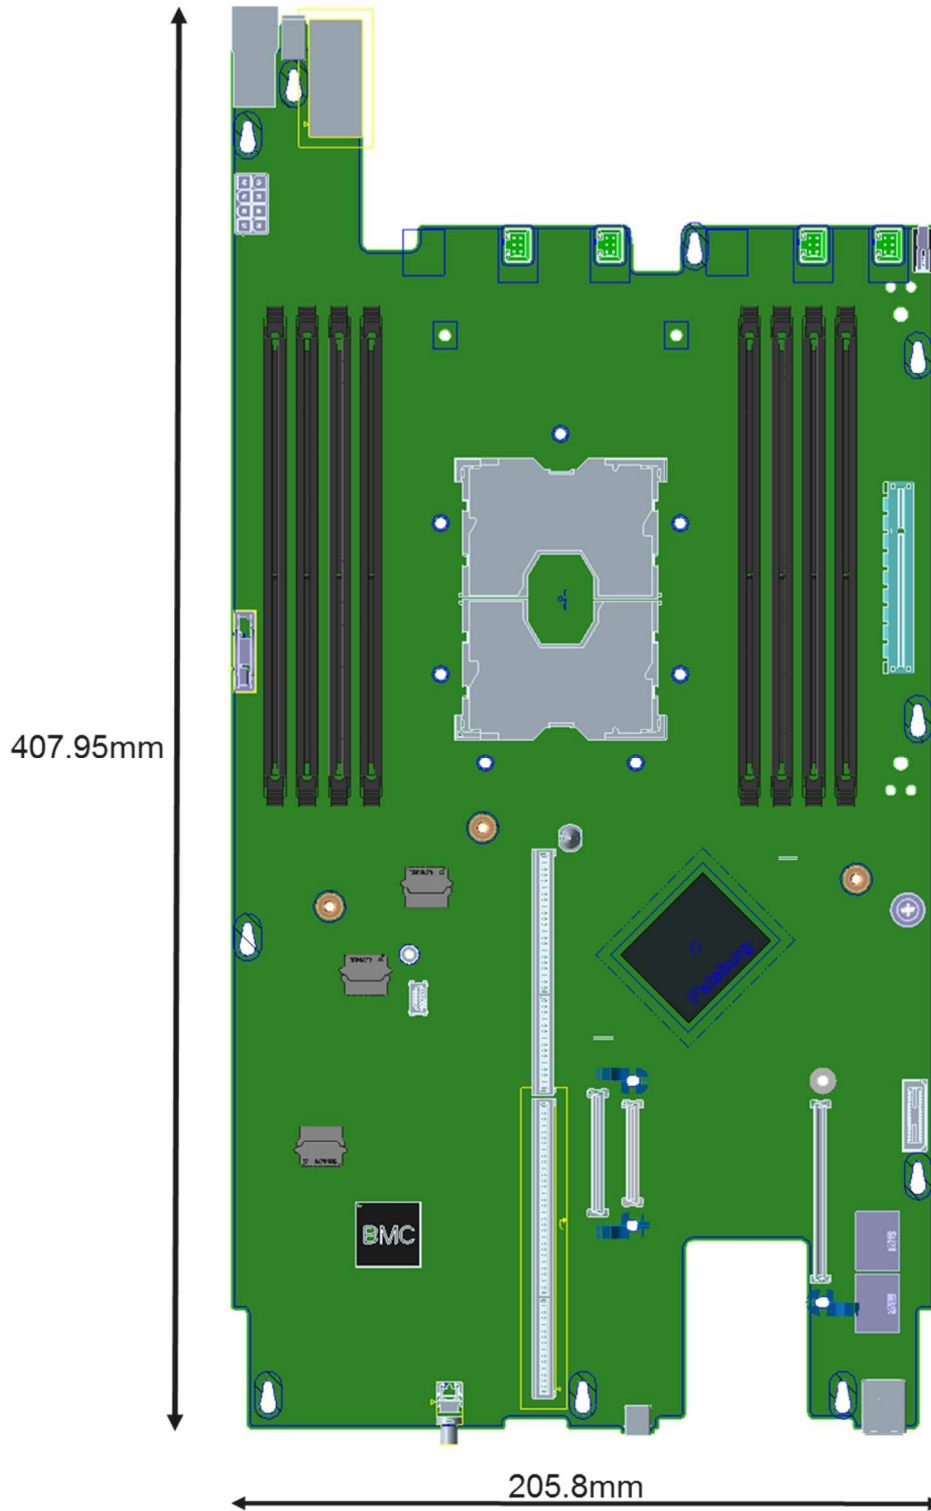


Figure 3: Server Main Board Dimension

3.2.3 Processor

The Cascade Lake SP is the next generation of 64-bit, multi-core server processor built on 14-nm process technology. The processor supports up to 46 bits of physical address space and 48 bits of virtual address space. The processor is designed for a platform consisting of at least one Cascade Lake SP processor and the Platform Controller Hub (PCH). Included in this family of processors are integrated memory controller (IMC) and an Integrated I/O (IIO) on a single silicon die.

All processor types support up to 48 lanes of PCI Express* 3.0 links capable of 8.0 GT/s, and 4 lanes of DMI3/PCI Express* 3.0. It features 2 Integrated Memory Controllers (IMC), each IMC supporting up to 3 channels of DDR4 DIMMs with up to 2 DIMM per channel.

3.2.4 PCH

The PCH's core name of Purley platform is named "Lewisburg". It is Intel C620 series chipset.

Lewisburg Key Features:

- ACPI Power Management Logic Support, Revision 4.0a
- PCI Express* Base Specification Revision 3.0
- Integrated Serial ATA host controller, supports data transfer rates of up to 6 Gb/s on all ports.
- xHCI USB controller with SuperSpeed USB 3.0 ports
- Direct Media Interface
- Serial Peripheral Interface
- Enhanced Serial Peripheral Interface
- Flexible I/O—Allows some high speed I/O signals to be configured as PCIE root ports, PCIE uplink for use with certain PCH SKUs, SATA (and sSATA), or USB 3.0.
- General Purpose Input Output (GPIO)
- Low Pin Count interface, interrupt controller, and timer functions
- System Management Bus Specification, Version 2.0
- Integrated Clock Controller / Real Time Clock Controller
- Intel® High Definition Audio and Intel® Smart Sound Technology
- Integrated 10/1 Gb Ethernet
- Integrated 10/100/1000 Gigabit Ethernet MAC

- Supports Intel® Rapid Storage Technology Enterprise
- Supports Intel® Active Management Technology and Server Platform Services
- Supports Intel® Virtualization Technology for Directed I/O
- Supports Intel® Trusted Execution Technology
- JTAG Boundary Scan support
- Intel® QuickAssist Technology
- Intel® Trace Hub for debug.
- Innovation Engine
- ADR Support

3.2.5 Memory

There are six memory channels A,B,C,D,E,F in the server board design and it supports eight DIMM slots. CH A,B,D and E are 1DPC. CH C and F are 2 DPC. CH C and F also supports Intel AEP. Here is the population matrix.

	Channel A	Channel B	Channel C		Channel D	Channel E	Channel F	
	Slot 0	Slot 0	Slot 0	Slot1	Slot 0	Slot 0	Slot 0	Slot1
Config1	DRAM	DRAM	DRAM	DRAM	DRAM	DRAM	DRAM	DRAM
Config2	DRAM	DRAM	DRAM	Empty	DRAM	DRAM	DRAM	Empty
Config3	DRAM	DRAM	DRAM	AEP	DRAM	DRAM	DRAM	AEP
Config4	DRAM	DRAM	AEP	Empty	DRAM	DRAM	AEP	Empty

The DIMM identifiers on the silkscreen on the board provide information about the channel to which they belong. For example, DIMM_A0 is the first slot on Channel A of processor.




Figure 4: Server Sled DDR Physical Slots

3.2.6 PCIE Lanes

CPU PCIE Port1

The PCIE x16 lanes between CPU, PCH and OCP mezzanine connector are in group of PCIE x 8 lanes. That is shown in below figure. It can be adjustable by the optional resistors. Based on PCH SKU selection, there are two configurations for the PCIE lanes.

Configuration 1:

PCH is C621.

CPU PCIE P1AB PCIE x 8 lanes are connecting to OCP mezzanine connector B.

CPU PCIE P1CD PCIE x 8 lanes are connecting to OCP mezzanine connector A.

Configuration 2

PCH is C627 and supports full bandwidth for QuickAssist.

CPU PCIE port1 x 16 lanes are connecting to PCH directly.

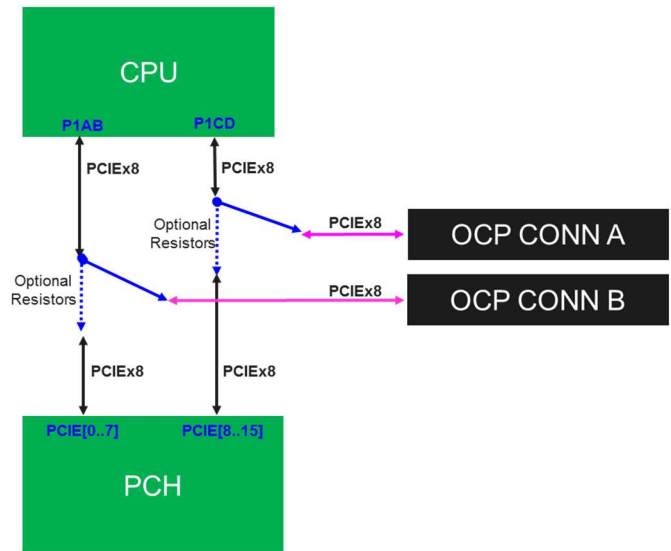


Figure 5: CPU PCIe Port1 Connection

CPU PCIe Port2

CPU PCIe P2AB PCIe x 8 lanes are connecting to 160 pin riser slot.
 CPU PCIe P2CD PCIe x 8 lanes are connecting to 200 pin riser slot.

CPU PCIe Port3

CPU PCIe P3ABCD PCIe x 16 lanes are connecting to 200 pin riser slot.

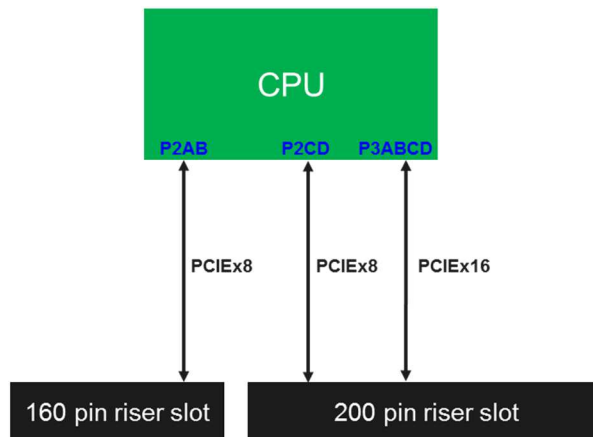


Figure 6: CPU PCIe Port2 and Port3 Connection

3.2.7 SATA

The Lewisburg PCH supports total 14 SATA-III 6Gbs ports. The PCH contains two SATA controller modes, while ACHI and Raid mode. Open Edge Server supports 8 SATAIII ports in SATA controller.

- Two SATA ports connect to 7 pins SATA connectors
- Four SATA ports connect to 160 pin riser slot.

- Two SATA ports connect to M.2 riser slot.

sSATA controller						SATA Controller							
SSATA 0	SSATA 1	SSATA 2	SSATA 3	SSATA 4	SSATA 5	SATA 0	SATA 1	SATA 2	SATA 3	SATA 4	SATA 5	SATA 6	SATA 7
						PCIE 0	PCIE 1	PCIE 2	PCIE 3	PCIE 4	PCIE 5	PCIE 6	PCIE 7

Figure 7: Purley SATA Ports Configuration

3.2.7.1 SATA PORT CONNECTIVITY




Figure 8: Server Sled SATA Diagram

3.2.8 USB

The Lewisburg PCH supports total 14 USB ports, 6 USB 2.0 ports and 8 USB 3.0 ports. The USB port distribution of Open Edge Server is as follows:

- ASPEED BMC AST2500 consumes 2 USB 2.0 ports (one 1.1 and one 2.0)
- 1 USB3.0 ports in front side
- 1 Mini USB port for debugging in front side

The USB ports on the products are not required to be powered from STBY.

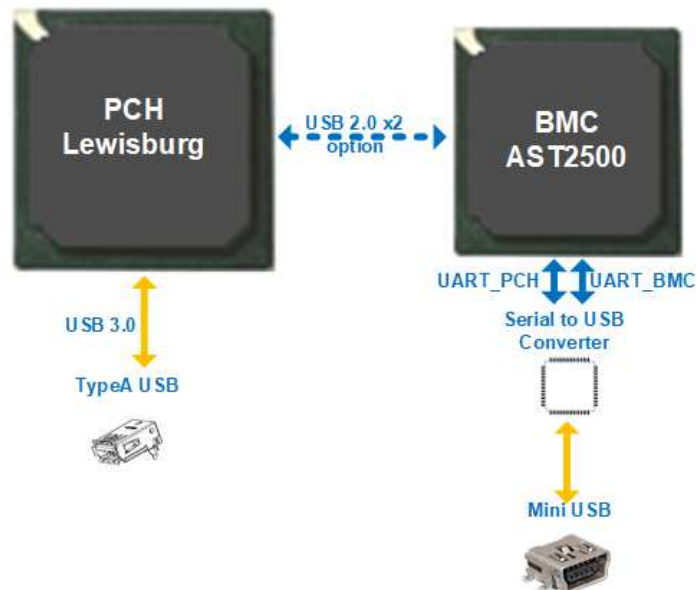


Figure 9: Server USB Ports

3.2.9 AST2500

The Server's Board Management Controller (**AST2500**) is a highly integrated single-chip solution, integrating several devices typically found on servers.

VGA Display Controller

- Fully IBM VGA compliant
- Maximum Display resolution: 1920x1200 32bpp@60Hz (reduced blanking)
- Support widescreen resolutions:
 - WXGA : 1280x800 32/16bpp @60Hz
 - WXGA+ : 1440x900 32/16bpp @60Hz
 - WSXGA+ : 1680x1050 32/16bpp @60Hz
 - FullHD+ : 1920x1080 32/16bpp @60Hz

DDR3L/DDR4 SDRAM Controller

- Support external 16-bit DDR3L/DDR4 SDRAM data bus width
- Maximum memory clock frequency
 - DDR3L: 800MHz (DDR3-1600)
 - DDR4: 800MHz (DDR4-1600)

GPIO Controller

- Directly connected to APB bus
- Support up to 228 GPIO pins, which are 29 sets
- Each GPIO sets can be programmed to accept command from ARM, LPC(SIO), or Coprocessor.
- Programmable output mode: Push-Pull or Open-Drain
- Some GPIOs support Schmitt type input buffer for noise immunity
- 4 out of the 228 GPIO pins are with 16mA driving strength, others are 8mA driving strength
- 16 out of the 228 GPIO pins that can support 1.8V mode.
- Support 8 sets of GPIO pass through (1 GPIO IN -- > 1 GPIO OUT) pin with internal switch control, it is useful for some button function control.

3.2.10 TPM

In Open Edge Server the PCH supports TPM specification 2.0 implemented with a TPM header on server sled. It supports SPI interface TPM 2.0 module.

3.2.11 UART

Open Edge Server provides Host and BMC UART interfaces for development purpose. Both UARTs connect to a UART to USB converter on the server board. The USB connect to a mini USB connector at the front IO. User can an USB cable to link the USB port to laptop to access the UARTs.

3.2.12 LOM

The LOM solution in Open Edge Server is Intel Ethernet controller I210-AT. I210-AT is a single port, compact, low power component that supports GbE designs. The I210 offers a fully-integrated GbE Media Access Control (MAC). The I210-AT enables 1000BASE-T implementations using an integrated PHY. Below are the features of I210-AT

- Operating Temperature: 0 to 70 °C.
- PCIE v2.1 (2.5 GT/s) x1 is used by the I210 as a host interface. In Open Edge Server, the PCIE interface connects to PCH PCIE port.
- Network Interface: 1000 Base-T. In server sled, the 1000 Base-T is routing on the server sled towards back plane




Figure 10: 1000 Base-T routing topology

- - It provides the below capabilities for the BMC on the server sled:
 1. A dedicated 1000 BASE-T Ethernet port for hardware management. User can remote manage the server sled via this interface.
 2. RMC can get the sled information via this interface.
 - Manageability: NCSI interface. This is accomplished by providing mechanisms by which manageability network traffic is routed to and from BMC.
 - Pass-Through (PT) is the term used when referring to the process of sending and receiving Ethernet traffic over the sideband interface. The I210 has ability to route Ethernet traffic to the host operating system as well as the ability to send Ethernet traffic over the sideband interface to an external BMC. The I210 supports two sideband interfaces:
 - SMBus
 - NC-SI
- Only one mode of sideband can be active at any given time.

In Open Edge Server, NCSI interface is connected to BMC. The usable bandwidth for either direction is up to 100 Mb/s for the NCSI interface.

3.2.13 Power connector for Accelerator Card

Open Edge Server can support FHFL DW accelerator card. The accelerator card needs the external 12V power feed via the cable. There is a 8 pin power connector on server sled (location J116) to offer the 12V power feed for accelerator. The power feed capacity of the power connector is 336W (max. current is 28A).

3.2.14 Thermal Design

3.2.14.1 Thermal solution for 1U sled

CPU Heat Sink

- Dimension=138*160*31 mm³ (+/-2 mm each)
- Material=Al base + Cu block + Cu Fin, w/RHE for saving TCO




Figure 11: 1U CPU Heatsink

Fan x4

- Dimension=4028 (single rotor)
- Voltage=12 V
- PWM Frequency= 25 KHz

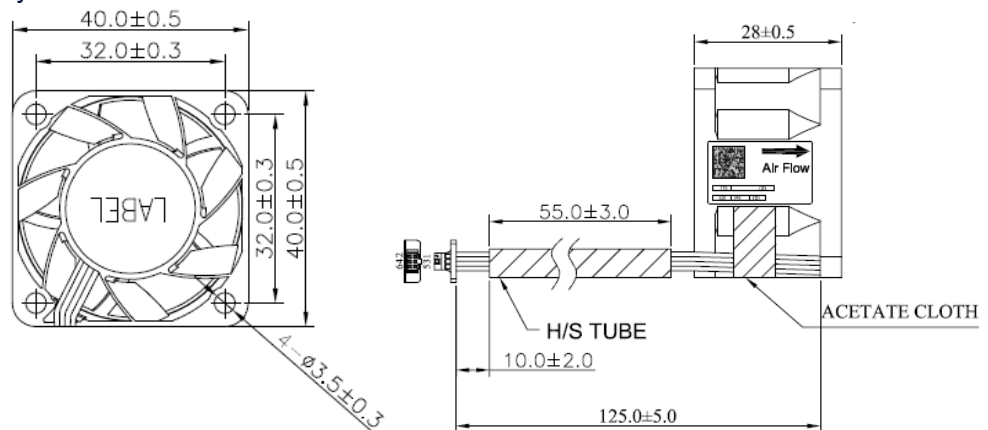


Figure 12: 4028 Fan Drawing

3.2.14.2 Thermal solution for 2U sled

CPU Heat Sink

- Dimension=135*160*63 mm³ (+/-2 mm each)
- Material=Al base + Cu block + Cu Fin, w/RHE for saving TCO




Figure 13: 2U CPU Heatsink

Fan x 2

- Dimension=8056 (dual rotors)
- Voltage=12 V
- PWM Frequency=25 KHz




Figure 14: 8056 Fan Drawing

4 Chassis Spec

4.1 3U System Overview

3U chassis can support below configurations:

- Server Sled
 - Five 1U server sleds at sled 1 to sled 5.
 - Two 2U server sleds at sled 2 and sled3.
- PSU
 - Support 1+1 redundant PSU
 - Support both AC/DC, DC/DC PSU.




Figure 15: 3U System Overview




Figure 16: 3U System Front View

4.1.1 3U Chassis Dimension

3U chassis dimension is 440mm x 430mm x 130.55mm (WxDxH)

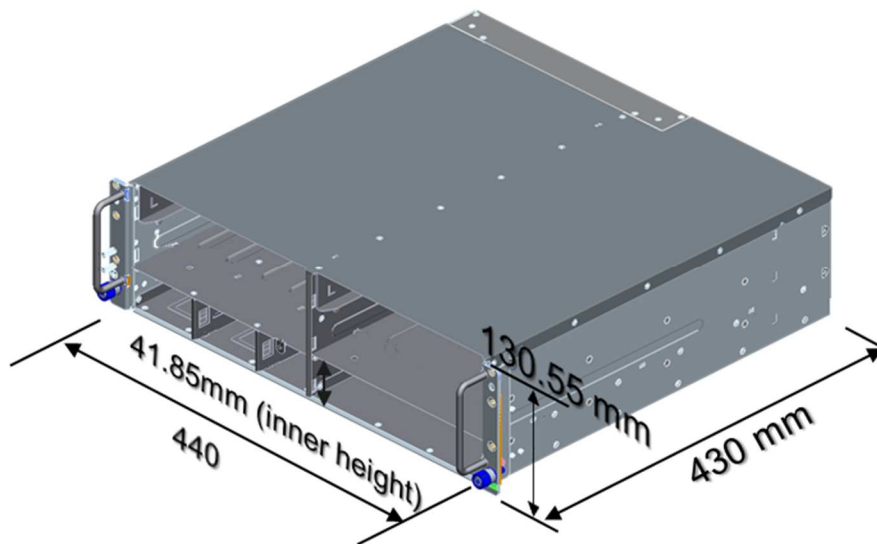


Figure 17: 3U Chassis Dimension

4.1.2 2U Sled Overview

Dimension 475 mm x 215 mm x 83.55 mm (DxWxH). It can support below configurations:

- Two M.2 cards
- One x16 FHH(/F)L PCIE card at PCIE slot3
- Two x8 FHHL PCIE card at PCIE slot1 and slot2
- 2 local 2,5" 9.5mm/7mm wide NVMe or SATA disks
- 2 local 2,5" 15mm wide NVMe or SATA disks




Figure 18: 2U- 3 Slot Sled Overview View

8056 Fan Module



Figure 19: 2U- 3 Slot Sled Top View

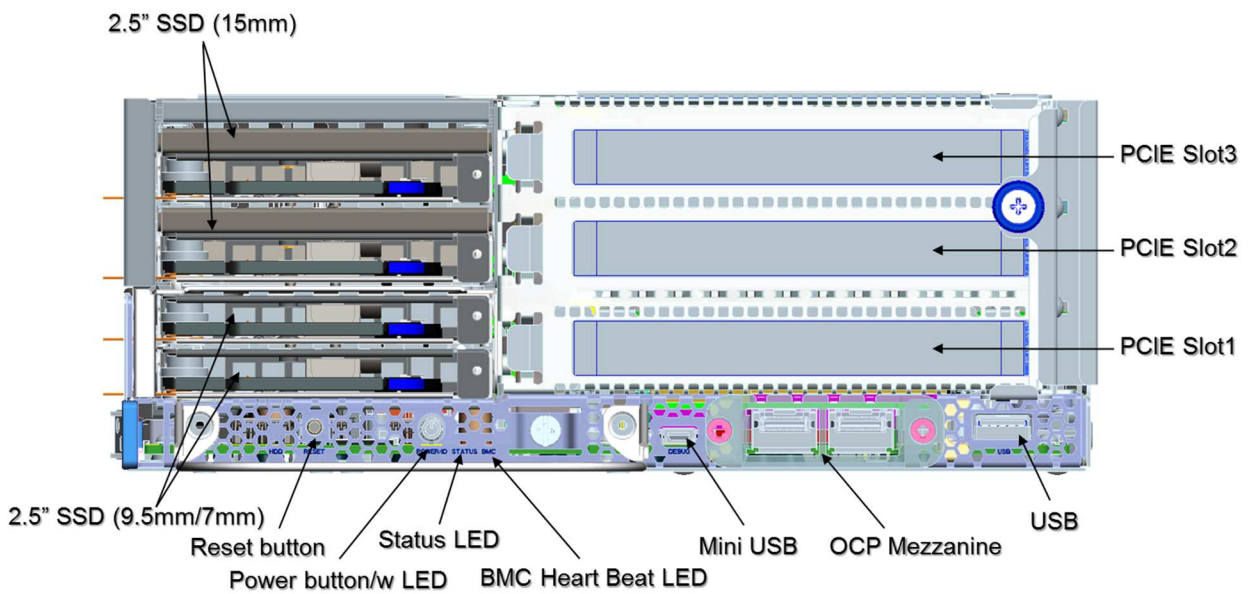


Figure 20: 2U- 3 Slot Sled Front View

4.1.3 1U Sled Overview

Dimension 475 mm x 215 mm x 41 mm (DxWxH). It can support below configurations:

- Two M.2 cards
- One x16 FHHL PCIE card
- 2 local 2,5" 9.5mm/7mm wide NVMe or SATA disks




Figure 21: 1U Sled Overview



Figure 22: 1U Sled Top view




Figure 23: 1U Sled Front view