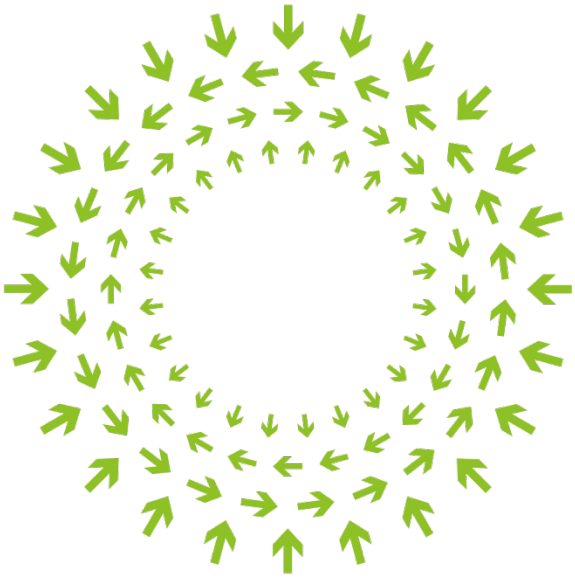

OPEN

Compute Project

Intel Motherboard Hardware v2.0

Author: Harry Li, Engineer

Author: Amir Michael, Engineer

1 Scope

This document defines the technical specifications for the Intel motherboard used in Open Compute Project servers.

2 Contents

3 Overview

When data center design and hardware design move in concert, they can improve efficiency and reduce power consumption. To this end, the Open Compute Project is a set of technologies that reduces energy consumption and cost, increases reliability and choice in the marketplace, and simplifies operations and maintenance. One key objective is openness—the project is starting with the opening of the specifications and mechanical designs for the major components of a data center, and the efficiency results achieved at facilities using Open Compute technologies.

One component of this project is a custom motherboard. This document describes the Open Compute Project Efficient Performance Intel motherboard, which supports up to 16 DIMMs. The motherboard is power-optimized and barebones, designed to provide the lowest capital and operating costs. Many features found in traditional motherboards have been removed from the design.

3.1 License

3.2 CAD Models

The following CAD files are incorporated by reference as if fully set forth in this specification:

4 Efficient Performance Motherboard Features

The Efficiency Performance motherboard, built with the Intel® Xeon® E5-2600 processor, was originally was code-named the Sandy Bridge-EP motherboard.

4.1 Block Diagram

Figure 1 illustrates the functional block diagram of the motherboard.

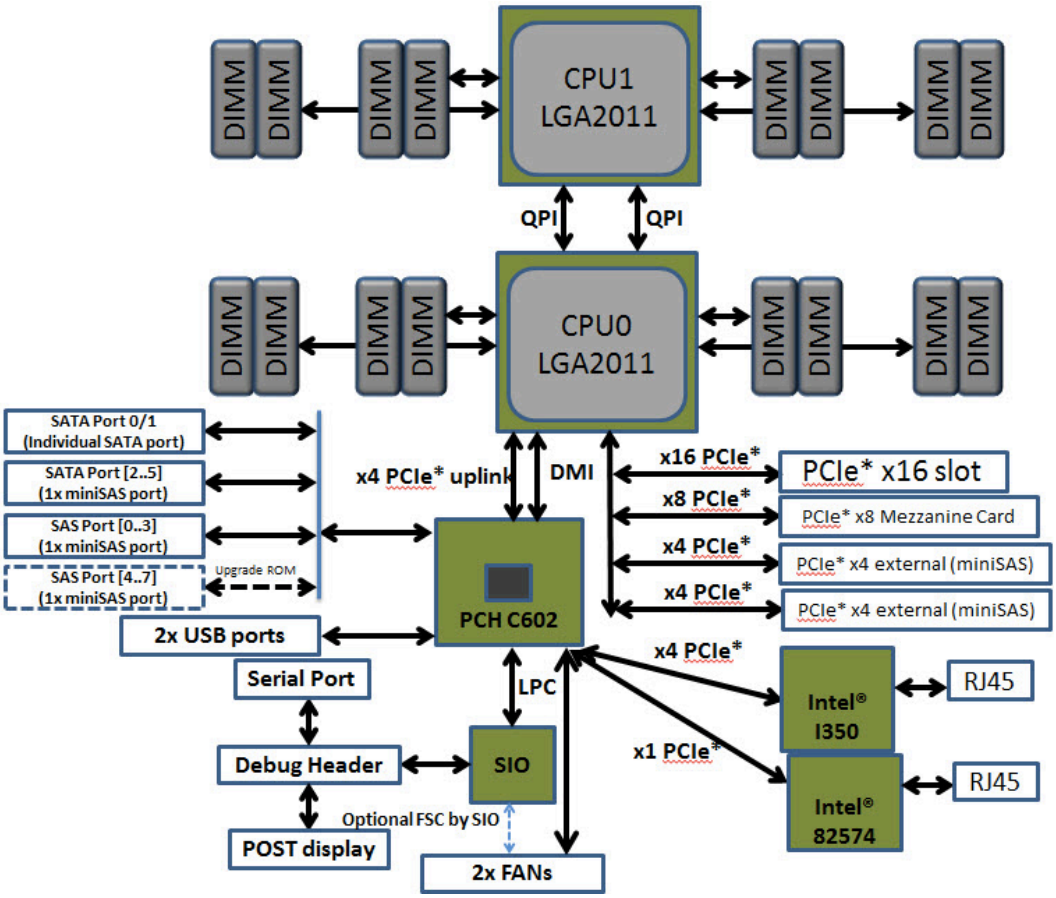

Figure 1 Efficient Performance Motherboard Functional Block Diagram

4.2 Placement and Form Factor

The motherboard's form factor is 6.5x20 inches. Figure 2 illustrates board placement. The placement shows the relative positions of key components, while exact dimension and position information is available in the mechanical DXF file. Form factor, PCIe* slot position, front IO port positions, PCIe* mezzanine card connector position, power connector, and mounting holes should be followed strictly, while other components can be shifted based on layout routing as long as relative position is maintained. As shown in Figure 28, one Open Compute chassis accommodates two motherboards. In order to remove and install one board without affecting the other board, the following internal connectors are placed as close as possible to front of the board in order to have easy frontal access:

- One vertical SATA signal connector and one SATA power connector.
- Debug card header.

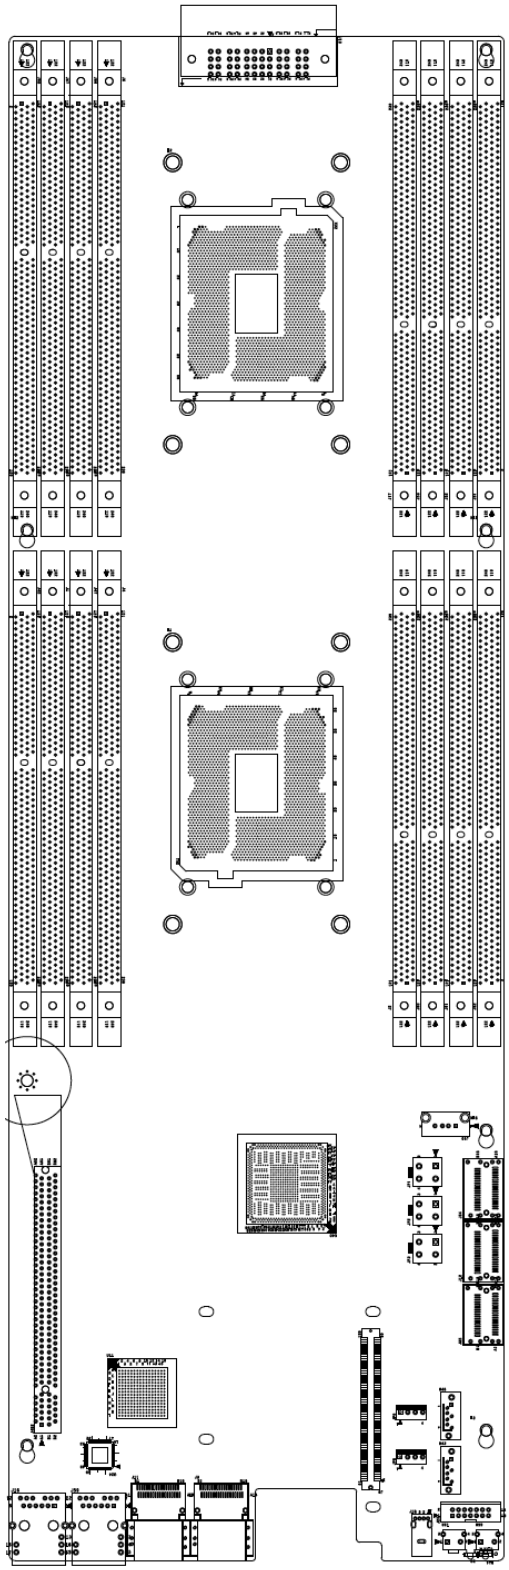

Figure 2 Efficient Performance Board Placement

4.3 CPU and Memory

The motherboard uses next generation Intel® Xeon® processor E5-2600 product family CPUs with a TDP (thermal design power) up to 115W. The motherboard supports these features:

- 2 Intel® Xeon® E5-2600 (LGA2011) series processors up to 115W
- 2 full-width Intel QuickPath interconnect (QPI) links up to 8 GT/s/direction
- Up to 8 cores per CPU (up to 16 threads with Hyper-Threading Technology)
- Up to 20 MB last level cache
- Single Processor Mode
- DDR3 direct attached memory support on cpu0 and cpu1 with:
  - 4 channel DDR3 registered memory interface on processors 0 and 1
  - 2 DDR3 slots per channel per processor (total of 16 DIMMs on the motherboard)
  - RDIMM/LV-RDIMM (1.5V/1.35V), LRDIMM and ECC UDIMM/LV-UDIMM(1.5V/1.35V)
  - Single, dual, and quad rank DIMMs
  - DDR3 speeds of 800/1066/1333/1600 MHz
  - Up to maximum 512 GB memory with 32GB RDIMM DIMMs

The memory support matrix is described in Figure 3.

		2 SPC	
		1 DPC	2 DPC
1.5V	SR/DR	1066, 1333, 1600	1066, 1333, 1600
	QR	1066	800
	LRDIMM	1066, 1333	1066, 1333
1.35V	SR/DR	1066, 1333	1066, 1333
	QR	800	800
	LRDIMM	1066	1066

Figure 3 Efficient Performance Board Memory Support Table

4.4 Platform Controller Hub

The Efficient Performance board uses the next generation Intel® platform controller hub (PCH) Intel® C602 chipset, which supports the following features:

- 14 USB 2.0 ports (2 designed in this board: one front connector, one vertical onboard)
- Serial-attached SCSI (SAS)/SATA ports:
  - 2 individual SATA 6Gps ports from SATA port 0/1
  - 1 mini-SAS port from SATA port 2/3/4/5
  - 1 mini-SAS port from SAS port 0/1/2/3
  - 1 mini-SAS port from SAS port 4/5/6/7, enabled by upgrade ROM (follow the Intel upgrade ROM design guide for implementation)
- SPI interface
- SMBUS interface (master and slave)
- SPS-DM management engine targets DCMI 1.5
  - Fan speed control (FSC) with optional system fan connection

- Temperature sensors read from external hardware monitor devices through I2C
- Remote access supported through SMLink0 connected to I350, 82574L, mezzanine LAN card and PCIe* LAN card. Depending on system configuration and SPS-DM configuration, only one NIC is enabled for ME shared NIC out-of-band access. Remote access features include serial over LAN, remote power control (power on, power off, power cycle), and a system event log.

4.5 PCIe* Usage

The Intel® Xeon® E5-2600 processor product family (LGA2011) can provide up to 40 PCIe* Express 3.0 lanes and the Intel® C602 chipset PCH can provide up to 8 PCIe* Express 2.0 lanes. The PCIe* lanes are distributed on the motherboard according to Figure 4.

The Efficient Performance board only supports NTB on PCIe* port3a/IOU1 in x4 or x8 configuration. In addition, PCIe* port3a/IOU1 should be routed to x4 PCIe* external connectors #1 and #2.

Device	Number of PCIe* Lanes
x4 PCIe* PCH uplink	4 (from E5-2600 port 1a/IOU2)
x16 PCIe* slot	16 (from E5-2600 port 2a/IOU0)
x4 PCIe* external connector #1 (miniSAS)	4 (from E5-2600 port3a/IOU1)
x4 PCIe* external connector #2 (miniSAS)	4 (from E5-2600 port3b/IOU1)
x8 PCIe* Mezzanine Card	8 (from E5-2600 port3c/IOU1)
Intel® I350 NIC	4 (from Intel® C602 chipset port 1)
Intel® 82574 NIC	1 (from Intel® C602 chipset port 7)
Total number of lanes	41

Figure 4 PCIe* Lane Usage

4.6 PCB Stackup

The Efficient Performance board's PCB stackup and impedance control are defined in the following tables.

Layer	Plane Description		Copper (oz)	Thickness (mil)	Dielectric (er)
		Solder Mask		0.5	3.8
L1	TOP	Signal	0.5+1.0	1.9	
		PrePreg		2.7	3.5
L2	GND1	Ground	1.0	1.3	
		Core		4.0	3.6
L3	IN1	Signal	1.0	1.3	
		PrePreg		7.7	4.0
L4	GND2	Ground	1.0	1.3	
		Core		4.0	3.6
L5	IN2	Signal	1.0	1.3	
		PrePreg		12.0	4.3
L6	VCC1	Power	2.0	2.6	

		Core		4.0	3.6
L7	VCC2	Power	2.0	2.6	
		PrePreg		12.0	4.3
L8	IN3	Signal	1.0	1.3	
		Core		4.0	3.6
L9	GND3	Ground	1.0	1.3	
		PrePreg		7.7	4.0
L10	IN4	Signal	1.0	1.3	
		Core		4.0	3.6
L11	GND4	Ground	1.0	1.3	
		PrePreg		2.7	3.5
L12	BOT	Signal	0.5+1.0	1.9	
		Solder Mask		0.5	3.8
		Total		85.2	Tolerance: +/-8mil

Figure 5 Efficient Performance Board PCB Stackup

Trace Width (mil)	Air Gap Spacing (mil)	Impedance Type	Layer	Impedance Target (ohm)	Tolerance (+/- %)
4.0		Single	1,12	50	15.0
6.5		Single	1,12	40	15.0
5.0		Single	1,12	43	15.0
3.9	4.1	Differential	1,12	83	17.5
3.8	8.2	Differential	1,12	95	17.5
5.0	7.0	Differential	1,12	85	17.5
4.1	13	Differential	1,12	100	10
4.0		Single	3,5,8,10	53	10.0
4.5		Single	3, 5,8,10	50	10.0
5.0		Single	3, 5,8,10	48	10.0
8.0		Single	3, 5,8,10	38	10.0
3.9	4.1	Differential	3, 5,8,10	85	12.0
4.4	9.1	Differential	3, 5,8,10	95	12.0
5.0	6.5	Differential	3, 5,8,10	85	12.0
4.1	13	Differential	3, 5,8,10	100	10.0

Figure 6 Efficient Performance Board PCB Impedance Control

5 BIOS

The ODM is responsible for supplying and customizing a BIOS for the motherboard. The specific BIOS requirements are outlined in this section.

5.1 BIOS Interface and Size

The BIOS chip uses the platform controller hub's SPI interface. The ODM is responsible for selecting a specific BIOS chip that meets the required functionality. 16MB of space should be allocated for both the BIOS and the PCH management engine.

5.2 BIOS Socket

A socket on the motherboard must be used to mount the BIOS chip to ensure that the BIOS chip can be replaced manually. The BIOS socket is easily accessible; other components on the motherboard do not interfere with the insertion or removal of the BIOS chip. A DIP-type BIOS chip and compatible socket are used for easy insertion and removal.

5.3 BIOS Source Code

The BIOS source code comes from AMI EFI. The ODM is responsible for maintaining the BIOS source code to make sure it has the latest version from AMI and Intel.

5.4 BIOS Configuration and Features

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

- Unused devices disabled, including PCIe* lanes, PCI lanes, USB ports, and SATA/SAS ports
- Tuning CPU/chipset settings to reach minimized power consumption and best performance
- SPEC_Power is used as guidance for ODM to validate BIOS tuning results

5.5 BIOS Setup Menu

The ODM must provide a BIOS specification, which includes the complete BIOS, setup menu, and default settings. The setup menu allows its options to be configured before the operating system loads. The configuration options available through the boot menu include the following:

- Settings to adjust memory speed, QPI speed, Speed-step/Turbo mode, and CPU Cx power state.
- Setting for power feature after AC failure, default is set to power on.
- Settings for console redirection:
  - **PCH ME Virtual COM port:** With baud rate 115200, no flow control, and terminal type VT100.
  - **SIO COM1:** With baud rate 57600, no flow control and terminal type VT100.
  - **Auto:** The PCH ME virtual COM port is enabled by default. The BIOS switches to SIO COM1 automatically, depending on hardware strapping. Default option is "Auto".
- Setting for fan speed control (for SIO FSC enabled board only).
- Setting for altitude of server deployment location.
- Hardware health monitoring display.
- Setting for watchdog timer. Default is enabled and timeout value is 15 minutes.

- Event log viewing and clearing.
- Setting for ECC error threshold, available settings are 1, 4, 10, and 1000.
- If a CMOS checksum error occurs (for example, caused by BIOS update), the BIOS loads the system default automatically after displaying a message in the console for 5 seconds and rebooting the system to apply the update without waiting for user input.
- Setting to disable all "wait for keyboard input to continue" types of features.

5.6 PXE Boot

The BIOS supports Intel PXE boot. When PXE booting, the system first attempts to boot from the first Ethernet interface (eth0). If a PXE boot on the first Ethernet interface fails, the BIOS attempts to PXE boot from the second Ethernet interface (eth1).

The default boot device priority is:

1. Network (search through all available network interfaces)
2. HDD
3. CDROM
4. Removable Device

5.7 Other Boot Options

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

5.8 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 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, retaining current BIOS settings
  - Reboot

Additionally, the update tools have the following capabilities:

- Update from the operating system over the LAN – the OS standard is CentOS v5.2
- 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 five minutes to complete
- Can be scripted and propagated to multiple machines

5.9 Event Log

As an alternative to the PCH management engine-provided event log, an event log may be implemented through SMBIOS.

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

The ODM must provide a system access interface and application software 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 and device ID.

5.9.1 Logged Errors

- **CPU/Memory errors:** Both correctable ECC and uncorrectable ECC errors should be logged into the event log. Error categories include DRAM, Link, and L3 cache.
- **QPI errors:** Any errors that have a status register should be logged into the event log. Fatal or non-fatal classification follows the chipset vendor's recommendation.
- **PCIe* errors:** Any errors that have a status register should be logged into the event log, including root complex, endpoint device, and any switch upstream/downstream ports if available. Link disable on errors should also be logged. Fatal, non-fatal, or correctable classification follows the chipset vendor's recommendation.
- **POST errors:** All POST errors detected by the BIOS during POST are logged into the event log.
- **Power errors:** Two power errors are logged:
  - 12.5V DC input power failure that causes all power rails on motherboard to lose power, including standby power.
  - Unexpected system shutdown during system S0/S1 while 12.5V DC input is still valid.
  -
- **MEMHOT# and PROCHOT# errors:** Memory hot errors and processor hot errors are logged. The error log identifies the error source as internally coming from the processor or memory, or externally coming from the voltage regulator.
- **Fan failures:** Fan failure errors are logged if the fan speed reading is out of the expected ranges between the lower and upper critical thresholds. The error log also identifies which fan fails.

5.9.2 Error Threshold Settings

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

- **Memory Correctable ECC:** The threshold value is 1000. When the threshold is reached, the BIOS should log the event including DIMM location information and output DIMM location code through the debug card.
- **QPI errors:** Follow the chipset vendor's suggestion.
- **PCIe* errors:** Follow the chipset vendor's suggestion.

5.9.3 BIOS Error Codes

MRC fatal error codes listed in following table should be enabled for POST code output. The major and minor codes alternately display.

Fatal Errors	Major Code	Minor Code	Error Description
ERR_NO_MEMORY	0E8h		
ERR_NO_MEMORY_MINOR_NO_MEMORY		01h	1. No memory was detected via SPD read. No warning log entries available. 2. Invalid configuration that causes no operable memory. Refer to warning log entries for details.
ERR_NO_MEMORY_MINOR_ALL_CH_DISAB LED		02h	Memory on all channels of all sockets are disabled due to hardware memtest error.
ERR_NO_MEMORY_MINOR_ALL_CH_DISAB LED_MIXED		03h	No memory installed. All channels are disabled.
ERR_LT_LOCK	0E9h		Memory is locked by LT, inaccessible.
ERR_DDR_INIT	0EAh		DDR3 training did complete successfully.
ERR_RD_DQ_DQS		01h	Error on read DQ/DQS init.
ERR_RC_EN		02h	Error on Receive Enable.
ERR_WR_LEVEL		03h	Error on Write Leveling.
ERR_WR_DQ_DQS		04h	Error on write DQ/DQS.
ERR_MEM_TEST	0EBh		Memory test failure.
ERR_MEM_TEST_MINOR_SOFTWARE		01h	Software memtest failure.
ERR_MEM_TEST_MINOR_HARDWARE		02h	Hardware memtest failure.
ERR_MEM_TEST_MINOR_LOCKSTEP_MODE		03h	Hardware memtest failure in Lockstep channel mode requiring a channel to be disabled. This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.
ERR_VENDOR_SPECIFIC	0ECh		Memory population violation
ERR_DIMM_COMPAT	0EDh		UDIMMs and RDIMMs are both present DIMM vendor-specific errors.
ERR_MIXED_MEM_TYPE		01h	Different DIMM types are detected installed in the system.
ERR_INVALID_POP		02h	Violation of population rules.
ERR_INVALID_POP_MINOR_QR_AND_3RD_SLOT		03h	The 3rd DIMM slot cannot be populated when quad rank (QR) DIMMs are installed.
ERR_INVALID_POP_MINOR_UDIMM_AND_ 3RD_SLOT		04h	UDIMMs and SODIMMs are not supported in the third DIMM slot.
ERR_INVALID_POP_MINOR_UNSUPPORTED _VOLTAGE		05h	Unsupported DIMM Voltage.
ERR_MRC_COMPATIBILITY	0EEh		Number of active Home Agent

			(detected) exceeds the pre-defined number of Home Agent.
ERR_MRC_STRUCT	0EFh		Indicates a CLTT table structure error. A DIMM is populated in the 3rd slot when quad rank DIMM is present in the channel.
ERR_INVALID_BOOT_MODE		01h	Boot mode is unknown.
ERR_INVALID_SUB_BOOT_MODE		02h	Sub boot mode is unknown.

Figure 7 BIOS Error Codes

6 Hardware Monitoring

The motherboard does not employ a traditional out of band monitoring solution. The ODM needs to provide a system access interface and application to retrieve hardware monitoring sensor readings. Lm_sensors is the preferred tool for hardware monitoring under Linux; the ODM ensures Lm_sensors works. The sensors to be read include voltage, temperature, and fan speed.

The NCT6681 serves as both the super IO (SIO) and hardware monitor.

6.1 Thermal Sensors

The motherboard has five thermal sensors:

- Two to monitor temperatures for CPU0 and CPU1, retrieved through the CPU's temperature sensor interface (PECI).
- DIMM temperatures of CPU0 DIMM group and CPU1 DIMM group, retrieved through the CPU PECI interface.
- PCH temperature, retrieved through Intel® C602 chipset internal DTS, through PCH SMLink1.
- Inlet temperature, retrieved through the thermistor, and located in the front of the motherboard.
- Outlet temperature, retrieved through the thermistor, and located in the rear of the motherboard.

The sensors should make sure that no CPU throttling is triggered due to thermal issues, under the following environmental conditions:

- Inlet temperature lower than 30C (including 30C), and 0 inch H2O pressure
- Inlet temperature higher than 30C but lower than 35C (including 35C), and 0.01 inch H2O pressure

The sensors should make sure that the total airflow rate for the chassis is lower than 89CFM, including PSU.

In the event that one fan fails, an inlet temperature of 30C with 0 inch H2O pressure environment is used to verify thermal sensors.

6.2 Fan Connection

The motherboard has fan tachometers and PWM connections to two system fans through the midplane. See section 8.2.

6.3 Fan Control Algorithm

The motherboard supports auto fan speed control for the system fans connected to it. The ODM must provide an optimized fan control algorithm based on the thermal solution of the system including fan, heat sink, and air duct. Fan speed control should set system fans running at lowest speed and provide enough damping to avoid speed vibration.

7 Platform Controller Hub Management Engine

The Intel® C602 chipset PCH management engine (ME) firmware (FW) is required for system operation. SPS-SiEn provides basic functionality required to have a fully functional platform. The ODM must incorporate the correct and latest SPS-SiEn based on the Intel released version for this motherboard. SPS-DM provides support for DCMI Rec 1.5. The ODM must also implement SPS-DM in the motherboard.

The OCP Intel motherboard is the first BMC-less platform for cloud computing that provides DCMI 1.5-based out-of-band management using the Intel PCH ME.

7.1 Fan Speed Control

The ODM must enable fan speed control (FSC) by SPS-DM as the default, with super IO (SIO) as an optional FSC. This includes hardware circuit and stuffing options. Temperature sensors and the FSC algorithm are the same as for SIO. The PCH ME can update the FSC configuration both locally (through the CentOS host OS) and remotely (through OOB); these updates take effect immediately without rebooting.

SPS-DM provides a fan PWM upper limit for fan boost during fan failure or ME firmware update.

7.2 Power and Thermal Monitoring

The ODM must implement SPS-DM to support platform power monitoring. To enable power monitoring/limiting for processors and memory, SPS-DNM is required (this is not Power On Record for the board). This is accessible both through in-band and out-of-band by LAN.

The ODM should implement SPS-DM to support thermal monitoring, including processor, memory, chipset, and inlet and outlet temperatures. Inlet and outlet temperature sensors must be compatible with Intel's ME/firmware supported device list. To ensure accuracy, use TI TMP421 to detect inlet and outlet temperatures.

7.3 Serial-Over-LAN (SOL)

The ODM should implement SPS-DM to support serial-over-LAN through the single shared network interface available on the motherboard. The shared network interface can come from LOM (82574 or I350) and NIC on the mezzanine card or the PCIe* card.

7.4 Remote Power Control

The ODM must implement SPS-DM to support remote system power on/off and reboot over LAN.

7.5 System Identification

Power LED blinking is used as system identification. The ODM should connect GPIO (GPIO30/MGPIO1) from the PCH ME to the power LED so the PCH ME can do system identification when it receives a chassis identification command through DCMI.

7.6 System Event Log

The ODM must implement SPS-DM to support a system event log. The BIOS can thus use the PCH ME SEL or SMBIOS for logging.

A digital sensor enables the PCH ME to log battery failure events, to report battery health status by the ME through MGPIO, and trigger event logging once the battery fails. The ODM must implement related hardware circuits for this. After battery replacement in the event of a failure, a system AC cycle with minimum 1 minute off time is required to re-enable the battery failure sensor.

Analog sensor monitoring thresholds are defined in Figure 8.

Sensor Name	Sensor Number	Lower Critical	Upper Critical
Outlet temp	0x01		75
Inlet temp	0x07		40
PCH temp	0x08		85
P0 therm margin	0x09		-2
P1 therm margin	0x0A		-2
P0 DIMM temp	0x11		82
P1 DIMM temp	0x12		82
HSC input power	0x29		350
HSC input voltage	0x2A	10.7	13.3
SYS FAN0	0x46	497	13987
SYS FAN1	0x47	497	13987

Figure 8 Analog Sensors

Sensor Name	Sensor #	Event Data0	Event Data1	Event Data2
CPU therm trip	0x10	[1]=1b, thermal trip		
Pwr thresh evt	0x3B	[1]=1b, limit exceeded		
Battery Mon	0x45	[0]=1b, low asserted		
SEL Status	0x5F	[1]=1b, SEL clear [8]=1b, SEL rollover		
DCMI Watchdog	0x60	[0]=1b, timer expired [1]=1b, hard reset [2]=1b, power down [3]=1b, power cycle [8]=1b, timer interrupt		
Processor Fail	0x65	[4]=1b, FRB3		

TSOD SMBus Sts	0x6E	0x00~0x03		
Chassis Pwr Sts	0x70	[0]=1b, power off [1]=1b, power cycle [4]=1b, AC lost		
Thermal Limit 1	0x7D	[1]=1b, limit exceeded		
SPS FW Health	0x17	[7,6]=10b [5,4]=10b [3..0]=00h, firmware status	=0, Forced GPIO recovery =1, Image execution failed =2, Flash erase error =3, Flash corrupted =4, Internal error = 5,BMC no respond =6, Direct Flash update by BIOS =7, Mfg. error =8, flash integrity error	<Extended error code>
POST Err	0x2B	[7:6]=10b or 11b [5:4]=10b [3:0]=00b, firmware error	ED1[7:6]= 10b, LSB of POST error code ED1[7:6]= 11b, per IPMI spec	MSB of POST error code
ProcHot Ext	0x3C	[7:6]=10b [5:4]=10b [3:0]=0Ah, processor thermal throttling offset	[7:2] Reserved [1:0] oh = native, 1h = external (VR)	[7:5] CPU/VR number [4:0] Reserved
MemHot Ext	0x3D	[7:6]=10b [5:4]=10b [3:0]=09h, memory thermal throttling offset	[7:2] Reserved [1:0] oh = native, 1h = external (VR)	[7:5] CPU/VR # [4:3] Channel # [2:0] DIMM # [4:0] Reserved for VR HOT
Machine Chk Err	0x40	[7:6]=10b [5:4]=10b [3:0]=0Bh, UnCorrectable or 0Ch, Correctable	Error code ID	[7:5] CPU # [4:3] Source 00b = QPI 01b = LLC [2:0], if QPI, (0b = QPI 0; 1b = QPI 1), if LLC, core #
PCIe* Err	0x41	[7:6]=10b [5:4]=10b	[2:0] Function # [7:3] Device #	[7:0] Bus #

		[3:0]=07h, correctable or 08h, uncorrectable or 0Ah, bus fatal		
Other IIO Err	0x43	[7:6]=10b [5:4]=10b [3:0]=00h, other IIO	Error code ID	[7:5] CPU # [4:3] Reserved [2:0] Source 000b = IRP0 001b = IRP1 010b = IIO-Core 011b = VT-d
Memory ECC Error	0x63	[7:6]=10b [5:4]=10b [3:0]=00h, correctable or 01h, uncorrectable	[1:0] Logical Rank [7:2] reserved	[7:5] CPU # [4:3] Channel # [2:0] DIMM #

Figure 9 Discrete and Event-Only Sensors

7.7 Firmware Update

The ODM must provide tool(s) to update the management engine firmware remotely, which does not require any physical input at the system. Remote update means either through out-of-band by the PCH ME or through logging into the local OS (CentOS 5.2) over the network. A remote ME firmware update may take a maximum of 5 minutes to complete and requires no more than one reset cycle to the system. The tool must support CentOS 5.2, should support updating the ME and BIOS together or separately, and must also provide an option to update only the operational ME region or the entire ME region.

8 Midplane

The midplane is a PCB board that functions as a bridge between the system fans, power supply (PSU), and both motherboards. Its form factor is 2x13 inches.

8.1 PSU Connector

The midplane has one FCI 51939-582 male right angle header, which is mated directly with the PSU for 12.5VDC input. Figure 10 shows the pin definition and direction based on the PSU.

Pin	Signal	Direction	Description	Usage
P1, P2	P12V	Power	12.5VDC	12.5VDC
P3, P4	GND	Power	Ground	Ground
A1	AUX_RTN_GND		Signal return	NC
A2	BACKUP_N	Output	PSU backup mode indication	NC
B1	SHARE_SEL_1	Input	PSU mode selection	NC
B2	SHARE_SEL_2	Input	PSU mode selection	NC
C1	GREEN_LED_N	Output	Low active	Connect to bi-color LED

C2	YELLOW_LED_N	Output	Low active	Connect to bi-color LED
D1	RED_LED_N	Output	Low active	Connect to LED
D2	P5V_AUX	Power	5V for LED, 50mA limited	LED power

Figure 10 Midplane to PSU Connector Pin Definition

For the PSU LED, the midplane provides a 4-pin vertically shrouded 2.54mm pitch header with latch. This allows an LED cable to extend the PSU LED to the chassis front. The PSU connector pins C1 and C2 connect to one bi-color (green/yellow) LED with a common anode. Pin D1 is connected to one red LED. Pin D2 is 5V and used for an LED anode. Both are 3mm LEDs. A current limit resistor is required for each LED signal.

Pin	Description
1	GREEN_LED_N
2	YELLOW_LED_N
3	RED_LED_N
4	P5V_AUX

Figure 11 PSU LED Header Pin Definition

When the PSU's red LED blinks (at 1Hz, 50% duty-cycle), it indicates a PSU fan failure.

8.2 Fan Connectors

The midplane has connectors for the four system fans. The connector signals comply both mechanically and electrically with the specifications defined in the *4-Wire Pulse Width Modulation (PWM) Controlled Fans Specification* Revision 1.3 September 2005 published by Intel Corporation. Each fan is driven by a dedicated PWM signal. Figure 12 defines the proper pinout of the connector.

Pin	Description
1	GND
2	12VDC
3	Sense
4	Control

Figure 12 Fan Header Pinout

A fan tachometer signal from each fan is routed to acquire fan speed. The midplane directly delivers 12.5V power to the fan connector. If one motherboard is not powered on, then its two corresponding fans are turned off to save power.

8.3 Motherboard Connectors

The midplane has two FCI 51770-044 female right-angle power/signal connectors (2P+16S+2P: 4 power blades and 16 signals). The motherboard -- with the mated FCI 51730-162 male right angle header -- slides in and mates with one of the FCI headers on the midplane. Figure 13 shows the pin definition of the 2P+16S+2P connector; the direction is based on the midplane.

Pin	Signals	Direction	Description
P1, P2	P12V	Power	12.5VDC
P3, P4	GND	Power	Ground
A1	SMB_ALT_N	Output	SMBUS alert signal from hot-swap controller
A2	TACH1A	Output	Reserved for extra fan tachometer on FAN1
A3	TACH2A	Output	Reserved for extra fan tachometer on FAN2
A4	RSVD		Reserved for future
B1	SCLK	Bi-direction	SMBUS CLOCK
B2	SDATA	Bi-direction	SMBUS DATA
B3	MB_ON	Input	Indicates that motherboard starts powered on
B4	PSU_PG	Output	Indicates that PSU 12.5VDC output is ready
C1	FAN1_TACH	Output	System fan #1 tachometer
C2	FAN1_PWM	Input	System fan #1 PWM
C3	FAN2_TACH	Output	System fan #2 tachometer
C4	FAN2_PWM	Input	System fan #2 PWM
D1 (short pin)	MATED_N	Input	Low active, indicates motherboard is fully mated
D2	MATED_GND_RTN		Connected to GND in midplane
D3	MB_ID	Output	Motherboard ID = 0 (left), 1 (right)
D4	FAN_FAIL_N	Output	PSU fan failure detected

Figure 13 Midplane to Motherboard Connector Pin Definition

8.4 Motherboard Power-up Delay

While running on AC power, in order to avoid both motherboards powering up at the same time and drawing larger than normal current, the midplane introduces a delay between the 12.5V power delivered to each of the two motherboards. The delay time can be set between 1 second and 1 minute, with 30 seconds as the default configuration.

The power-up delay behaves as follows:

- When both MB0 and MB1 are installed and AC power is applied, MB0 powers on first, then after 30 seconds (the timer delay), MB1 powers on.
- When both MB0 and MB1 are operating, and you remove and re-insert a motherboard, there is no delay for it to power on again.
- When both MB0 and MB1 are operating, and you remove and re-insert both boards, a 30 second timer delay still occurs between MB0 and MB1 powering on.
- When only MB0 is installed and AC power is applied, there is no delay when it powers on.
- When only MB1 is installed and AC power is applied, there is no delay when it powers on.
- With one motherboard is operating, and another motherboard is inserted, there is no delay when it powers on.

- If no motherboards are installed and AC power is applied, then both MB0 and MB1 are inserted, there is a 30 second timer delay between MB0 and MB1 powering on.

8.5 Hot Swap Controller

In order to have better control of the 12.5VDC power input to each motherboard, the ODM should include two hot swap controllers (ADI ADM1276, one for each motherboard) on the midplane. The hot swap controller provides:

- Inrush current control when the board is inserted and the server is powered on.
- Current limiting protection for short circuit. Overcurrent trip point should be set to 33A with a jumper setting adjustable to 40A.
- PMBUS interface to enable the PCH ME to report server input power.
  - Report server input power and log event if it triggers upper critical threshold.
  - Report input voltage (up to 1 decimal point) and log event if it triggers either lower or upper critical threshold.
  - Log status event based on hot swap controller's status register.

Current sense resistor value for hot swap controller should be set to 0.5 mohm. Both hot swap controllers on midplane should have their SMBUS address set to 0x20 (7-bit format).

9 Power System

9.1 Input Voltage

9.1.1 Input Voltage Level

The nominal input voltage delivered by the power supply is 12.5VDC. The motherboard can accept and operate normally with an input voltage tolerance range between 10.8V and 13.2V. The motherboard's undervoltage protection level is 10V or less.

9.1.2 Capacitive Load

To ensure compatibility with the system power supply, the motherboard cannot have a capacitive load greater than 4000µF. The capacitive load of the motherboard cannot exceed the maximum value of 4000µF under any operating condition listed in section 12, which defines environmental conditions.

9.1.3 Input Connector

The power input connector is an FCI 51733-009LF right-angle press-fit header.

9.2 CPU Voltage Regulation Module (VRM)

9.2.1 CPU Maximum Power

The motherboard can handle a processor with a maximum thermal design power (TDP) of 115W for the Efficient Performance board.

9.2.2 CPU VRM Optimizations

The CPU VRM is optimized to reduce cost and increase the efficiency of the power conversion system. The ODM should use only the minimum number of required phases

to support the maximum CPU power defined in 9.2.1. A PSI (power state indicator) allows the shedding of unused phases, letting the VRM operate at its peak efficiency.

9.2.3 CPU VRM Efficiency

The minimum efficiency for the CPU VRM is 91% over the 30% to 90% load range and 93% over the 50% to 70% load range, measured from the 12.5V input to the VRM output.

9.3 Hard Drive Power

The motherboard supplies power to the system's 14 hard drives. The drives require 12VDC and 5VDC power sources. Power is delivered through a traditional 4-pin floppy disk power connector described in Figure 14.

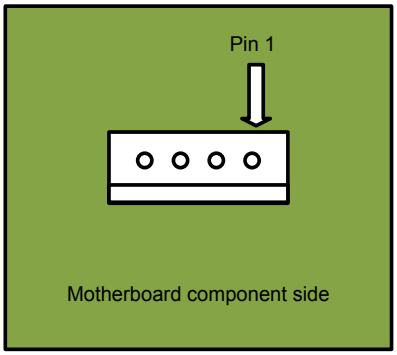

Pin	Description
1	+5VDC
2	GND
3	GND
4	+12VDC

Figure 14 Drive Power Connector

For SATA ports inside the miniSAS connector, power will be delivered through a 4-pin (2x2) ATX power connector, which fans out into 4 standard SATA power cables. Pin definition is described in Figure 15.

Pin	Description
1	GND
2	GND
3	+5VDC
4	+12VDC

Figure 15 4 Pin ATX Power Connector

9.3.1 Power Requirements

In order for the motherboard to supply 12.5VDC power to the hard drives, the PCB traces must support 14A of continuous power (1A per drive) on the 12.5VDC power rail. In order for the system's 5VDC to supply power to the hard drives, its regulator must support an additional 10.5A (0.75A per drive) of continuous power on the 5VDC power rail. The motherboard must support the inrush current required to start each drive from idle.

9.3.2 Output Protection

The 5V disk output power regulator protects against shorts and overload conditions.

9.3.3 Spin-up Delay

When a hard drive spins up after the system powers on, it draws excessive current on both the 12V and 5V rails. The peak current may reach the 1.5A-2A range in 12V. Each of the 14 hard drives must spin up in sequence. The BIOS implements a 5 second delay between each hard drive spinning up. To enable the hard drive's spin-up delay function, set pin 11 of the SATA hard drive's power cable to NC (No Connection).

9.4 System VRM Efficiency

The ODM supplies high efficiency VRMs for all other voltage regulators over 20W not defined in this specification. All voltage regulation modules over 20W have 91% efficiency over the 30% to 90% load range.

9.5 Power On

The motherboard powers on upon application of power to the input connector. The use of a power button is not required. The motherboard always resumes operation upon restoration of power in a power failure event.

9.6 VRM Design Guidelines

The VRM OCP setting should be set to double the maximum loading current, under conditions that it should meet the specification and ratings of the components used in the VRM. For a VRM that requires firmware or a power code or configuration file, the ODM should maintain version control to track all the releases and changes between each version, and provide a method to retrieve the version through application software during system runtime. This software runs under CentOS.

10 I/O System

This section describes the motherboard's I/O features.

10.1 PCIe* x16 Slot/Riser Card

The Efficient Performance motherboard has one PCIe* x16 slot, which holds an x16 PCIe* Gen 3 signal from the CPU. The slot location and detailed dimensions are described in the mechanical DXF file. The motherboard also has a PCIe* riser card so two full-height PCIe* cards can be inserted horizontally and locked in position. Its form factor is 2x4.66 inches.

Figure 16 PCIe* Riser Card

The reserved pins on the PCIe* x16 slot on the motherboard are described in Figure 17.

Pin	Pin Defined	Description
A7	LAN_SMB_CLK	SMBUS CLOCK from SMLINK0 of PCH
A8	LAN_SMB_DAT	SMBUS DATA from SMLINK0 of PCH
B12	LAN_SMB_ALERT_N	SMBUS Alert signal to SMLINK0 of PCH
A32	CLK_100M_P	Extra 100MHz clock for second PCIe* slot on riser card
A33	CLK_100M_N	
A50	SLOT0_CONFIG	Lower slot on riser card has 1x8 (high), 2 x4 (low)
B82	SLOT1_CONFIG	Higher slot on riser card has 1x8 (high), 2 x4 (low)
B17	SLOT1_CPRSNT1_N	CPRSNT1# for SLOT1 on PCIe* riser card
B31	SLOT1_CPRSNT2_N	CPRSNT2# for SLOT1 on PCIe* riser card
B48	SLOT0_CPRSNT1_N	CPRSNT1# for SLOT0 on PCIe* riser card
B81	SLOT0_CPRSNT2_N	CPRSNT2# for SLOT0 on PCIe* riser card

Figure 17 PCIe* x16 Slot Reserved Pin Usage on Motherboard

The reserved pins on the PCIe* x16 slot 0 (low) on the riser card are described in Figure 18.

Pin	Pin Defined	Description
A32	LAN_SMB_CLK	SMBUS clock from SMLINK0 of PCH
A33	LAN_SMB_DAT	SMBUS data from SMLINK0 of PCH
A50	LAN_SMB_ALERT_N	SMBUS alert signal to SMLINK0 of PCH
B48	SLOT0_CPRSNT1_N	CPRSNT1# for SLOT0 on PCIe* riser card
B81	SLOT0_CPRSNT2_N	CPRSNT2# for SLOT0 on PCIe* riser card

Figure 18 PCIe* x16 Slot 0 (Low) Reserved Pin Usage on Riser Card

The reserved pins on the PCIe* x16 slot 1 (high) on riser card are described in Figure 19.

Pin	Pin Defined	Description
A32	LAN_SMB_CLK	SMBUS clock from SMLINK0 of PCH
A33	LAN_SMB_DAT	SMBUS data from SMLINK0 of PCH
A50	LAN_SMB_ALERT_N	SMBUS alert signal to SMLINK0 of PCH
B17	SLOT1_CPRSNT1_N	CPRSNT1# for SLOT1 on PCIe* riser card
B31	SLOT1_CPRSNT2_N	CPRSNT2# for SLOT1 on PCIe* riser card

Figure 19 PCIe* x16 Slot 1 (High) Reserved Pin Usage on Riser Card

To support OOB LAN access on the platform controller hub's management engine, a customized PCIe* card is needed to use these redefined reserved pins.

10.2 PCIe* External Connector

The Efficient Performance motherboard has two PCIe* x4 external connectors on the motherboard. These connectors can be used to build a PCIe* connection between two systems.

The PCIe* x4 connector can be hot inserted and removed. A PCIe* 3.0 re-driver is used for PCIe* external links and supports a miniSAS cable up to 2 meters long.

The connector is a miniSAS-4i right-angle connector. External PCI Express target device is TBD. Figure 20 shows the external PCIe* pin assignments. The design follows the *PCI Express External Cabling 1.0 Specification* (http://www.pcisig.com/members/downloads/specifications/PCIe*xpress/PCI_Express_External_Cabling_Rev1.0_updated.pdf).

Pin Numbers	Signals	Description
A2/A3, A5/A6, A13/A14, A16/A17	PER{0..3}{P/N}	Differential PCI Express receiver lanes
A1, A4, A7, A12, A15, A18	GND	Ground reference for Differential PCI Express lanes
A8	CPRSNT#	Cable installed/downstream subsystem powered up
A9	CPWRON	Upstream subsystem's power valid notification
A10	CWAKE#	Power management signal for wakeup events (optional)
A11	CPERST#	Cable PERST#
B2/B3, B5/B6, B13/B14, B16/B17	PET{0..3}{P/N}	Differential PCI Express transmitter lanes
B1, B4, B7, B12, B15, B18	GND	Ground reference for Differential PCI Express lanes
B8	SCLK/TX	SMBUS (PCH SMLINK0) CLOCK (optional UART TX from SIO)
B9	SDATA/RX	SMBUS (PCH SMLINK0) DATA (optional UART RX from SIO)
B10	3.3V/SYS_RST#	3.3V standby with 0 ohm in series (Reset signal to trigger system reset)
B11	SB_RTN	Signal return for single-ended sideband signals

Figure 20 External PCIe* Pin Assignments

10.3 PCIe* Mezzanine Card

The motherboard has one PCIe* x8 mezzanine card connector that holds the x8 PCIe* Gen 3 signal from cpu0 on the Efficient Peformance motherboard. The mezzanine card has two PCIe* x4 external connectors (miniSAS) and one mSATA miniPCIe* connector.

Pin Name	Pin	Pin	Pin Name	Pin Name	Pin	Pin	Pin Name
P12V	61	1	MEZZ_PRSNT1_N	GND	91	31	MEZZ_RX_DN<0>
P12V	62	2	P5V_AUX	MEZZ_TX_DP_C<1>	92	32	GND
P12V	63	3	P5V_AUX	MEZZ_TX_DN_C<1>	93	33	GND
GND	64	4	P5V_AUX	GND	94	34	MEZZ_RX_DP<1>
GND	65	5	GND	GND	95	35	MEZZ_RX_DN<1>
P3V3_AUX	66	6	GND	MEZZ_TX_DP_C<2>	96	36	GND
GND	67	7	P3V3_AUX	MEZZ_TX_DN_C<2>	97	37	GND
GND	68	8	GND	GND	98	38	MEZZ_RX_DP<2>
P3V3	69	9	GND	GND	99	39	MEZZ_RX_DN<2>
P3V3	70	10	P3V3	MEZZ_TX_DP_C<3>	100	40	GND
P3V3	71	11	P3V3	MEZZ_TX_DN_C<3>	101	41	GND
P3V3	72	12	P3V3	GND	102	42	MEZZ_RX_DP<3>
GND	73	13	P3V3	GND	103	43	MEZZ_RX_DN<3>
LAN_3V3STB_ALERT_N	74	14	MEZZ_CPRSNT1_N	MEZZ_TX_DP_C<4>	104	44	GND
SMB_LAN_3V3STB_CLK	75	15	MEZZ_CPRSNT2_N	MEZZ_TX_DN_C<4>	105	45	GND
SMB_LAN_3V3STB_DAT	76	16	SSD_PRSNT_N	GND	106	46	MEZZ_RX_DP<4>
PCIe*_WAKE_N	77	17	RST_PLT_MEZZ_N	GND	107	47	MEZZ_RX_DN<4>
DA_DSS	78	18	MEZZ_SMCLK	MEZZ_TX_DP_C<5>	108	48	GND
GND	79	19	MEZZ_SMDATA	MEZZ_TX_DN_C<5>	109	49	GND
SATA_TX+	80	20	GND	GND	110	50	MEZZ_RX_DP<5>
SATA_TX-	81	21	GND	GND	111	51	MEZZ_RX_DN

							<5>
GND	82	22	SATA_RX+	MEZZ_TX_DP_C<6>	112	52	GND
GND	83	23	SATA_RX-	MEZZ_TX_DN_C<6>	113	53	GND
CLK_100M_MEZZ2_DP	84	24	GND	GND	114	54	MEZZ_RX_DP<6>
CLK_100M_MEZZ2_DN	85	25	GND	GND	115	55	MEZZ_RX_DN<6>
GND	86	26	CLK_100M_MEZZ1_DP	MEZZ_TX_DP_C<7>	116	56	GND
GND	87	27	CLK_100M_MEZZ1_DN	MEZZ_TX_DN_C<7>	117	57	GND
MEZZ_TX_DP_C<0>	88	28	GND	GND	118	58	MEZZ_RX_DP<7>
MEZZ_TX_DN_C<0>	89	29	GND	GND	119	59	MEZZ_RX_DN<7>
GND	90	30	MEZZ_RX_DP<0>	MEZZ_PRSNT2_N	120	60	GND

Figure 21 PCIe* Mezzanine Card Connector Pin Definition

10.4 DIMM Connector

The motherboard uses a 30u" gold contact for the DDR3 DIMM connector.

10.5 Network

The motherboard has an Intel® 82574L Ethernet interface to the front RJ45 connector. It has a PCIe* x1 lane routed to the PCH.

The motherboard has an Intel® I350 dual port network chip. It has a single Ethernet interface to the front RJ45 connector. It has PCIe* x4 lanes routed to the PCH on the Efficient Performance board.

The BIOS supports PXE boot on all RJ45 ports on the motherboard.

Each RJ45 connector has two built-in LEDs. While facing the RJ45 connector, the left LED is green single color; solid on means the link is active and blinking means activity. The right LED is green/yellow dual color; green means 100M link speed while yellow means 1000M link speed.

10.5.1 Reboot on WOL in S0 State

Reboot on WOL (ROW) is a feature that repurposes the traditional Wake on LAN (WOL) signal to reboot the motherboard. While the system is in S0 state (running), when a WOL packet is received by the NIC, the wakeup signal generated by the NIC causes a hardware reboot of the motherboard. This is accomplished by tying the WOL interrupt pin of the NIC to the system's master reboot signal. ROW does not require the power supply to cycle its output.

There is an optional ROW connection for the WAKE# signal from PCIe* slot and external PCIe* connector, which gives optional ROW support for add-in cards and external PCIe* devices.

ROW is enabled by the NIC EEPROM, so the appropriate NIC EEPROM for the Intel® 82574 and Intel® I350 interface must be used. The motherboard also supports ROW on both the PCIe* LAN card and the mezzanine LAN card, which includes hardware circuit support and NIC EEPROM enabling.

10.5.2 Out of Band Network Access

The motherboard supports out of band (OOB) network access to the platform controller hub's management engine through network interfaces on the Intel® 82574, the Intel® I350, PCIe* LAN, and mezzanine LAN card interfaces. This includes all remote access features described in this specification. Only one network interface with OOB support is needed for each system configuration. For PCIe* LAN cards and mezzanine LAN cards, SMBUS and SMBAlert connections to PCH ME SMLINK0 are required.

10.6 USB Interfaces

The motherboard has two external USB ports located in the front of the motherboard. The BIOS supports the following USB devices:

- Keyboard and mouse
- USB flash drive (bootable)
- USB hard drive (bootable)
- USB optical drive (bootable)

10.7 SATA

The motherboard has an Intel® C602 chipset platform controller hub on board. Two single SATA ports from Intel® C602 chipset SATA port 0 and port 1 are connected to two single SATA connectors, which support SATA 6Gbps. Four SATA ports from Intel® C602 chipset SATA port 2 to 5 are connected to one miniSAS connector, which support SATA 3Gbps. Four SAS ports from Intel® C602 chipset SAS ports 0 to 3 are connected to one miniSAS connector, which support SATA 6Gbps. When the upgrade ROM is installed, the Intel® C602 chipset will support four extra SAS ports (4 to 7) that are connected to one miniSAS connector.

The Efficient Performance board fully supports all 8 SAS ports. The HDDs attached to all the SATA connectors follow the spin-up delay described in section 9.3.3.

10.8 Debug Header

The motherboard includes a debug header on the front of the motherboard to display POST codes (see 10.8.1). The debug header supports hot plugging.

The debug card has two 7-segment LED displays, one RS-232 serial connector, and one reset switch. The RS-232 serial port provides console redirection. The two 7-segment LED displays show BIOS POST code and DIMM error information. The reset switch triggers a system reset when pressed.

The connector for the debug header is a 14-pin, shrouded, vertical, 2mm pitch connector. Figure 22 is an illustration of the headers. The debug card has a key to match with the notch to avoid pin shift when plugging it in.

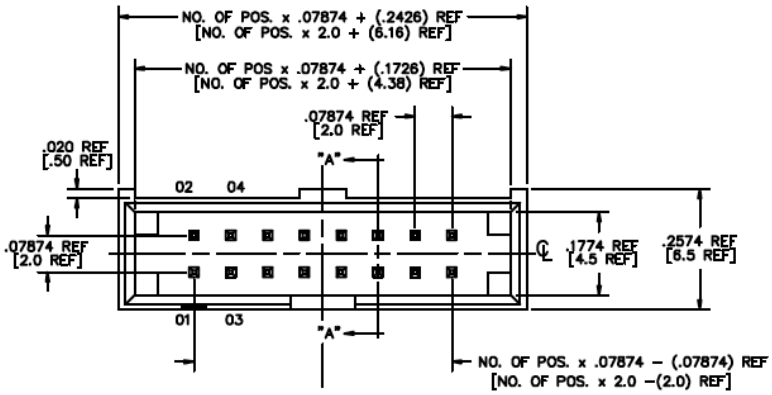

Figure 22 Debug Header

Pin (CKT)	Function
1	Low HEX character [0] least significant bit
2	Low HEX character [1]
3	Low HEX character [2]
4	Low HEX character [3] most significant bit
5	High HEX character [0] least significant bit
6	High HEX character [1]
7	High HEX character [2]
8	High HEX character [3] most significant bit
9	Serial transmit (motherboard transmit)
10	Serial receive (motherboard receive)
11	System reset
12	Serial console select (1=SOL; 0=local)
13	GND
14	VCC (+5VDC)

Figure 23 Debug Header Pin Definitions

10.8.1 Post Codes

POST codes are sent to the debug header in hexadecimal format via two hex codes. The hex codes can be driven by either the legacy parallel port (port 80) on the SIO, or 8 GPIO pins.

During POST, the BIOS should also output POST codes to the PCH ME via serial over LAN (SOL). When a SOL session is available during POST, the remote console should display POST codes.

During the boot sequence, the BIOS initializes and tests each DIMM. If a module fails initialization or does not pass the BIOS test, one of the following POST codes will flash on the debug card to 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 BIOS flashes the corresponding hex code indefinitely to allow time for a technician to service the system. The DIMM number count starts from the DIMM furthest from the CPU.

CPU	Code	Result
CPU0 (Channels 0 & 1)	A0	CPU 0 channel 0 DIMM 0 (furthest) failure
	A1	CPU 0 channel 0 DIMM 1 failure
	A2	CPU 0 channel 1 DIMM 0 failure
	A3	CPU 0 channel 1 DIMM 1 (closest) failure
CPU0 (Channels 2 & 3)	A4	CPU 0 channel 2 DIMM 0 (furthest) failure
	A5	CPU 0 channel 2 DIMM 1 failure
	A6	CPU 0 channel 3 DIMM 0 failure
	A7	CPU 0 channel 3 DIMM 1 (closest) failure
CPU1 (Channels 0 & 1)	B0	CPU 1 channel 0 DIMM 0 (furthest) failure
	B1	CPU 1 channel 0 DIMM 1 failure
	B2	CPU 1 channel 1 DIMM 0 failure
	B3	CPU 1 channel 1 DIMM 1 (closest) failure
CPU1 (Channel 2 & 3)	B4	CPU 1 channel 2 DIMM 0 (furthest) failure
	B5	CPU 1 channel 2 DIMM 1 failure
	B6	CPU 1 channel 3 DIMM 0 failure
	B7	CPU 1 channel 3 DIMM 1 (closest) failure

Figure 24 Efficient Performance Motherboard DIMM Error Code Table

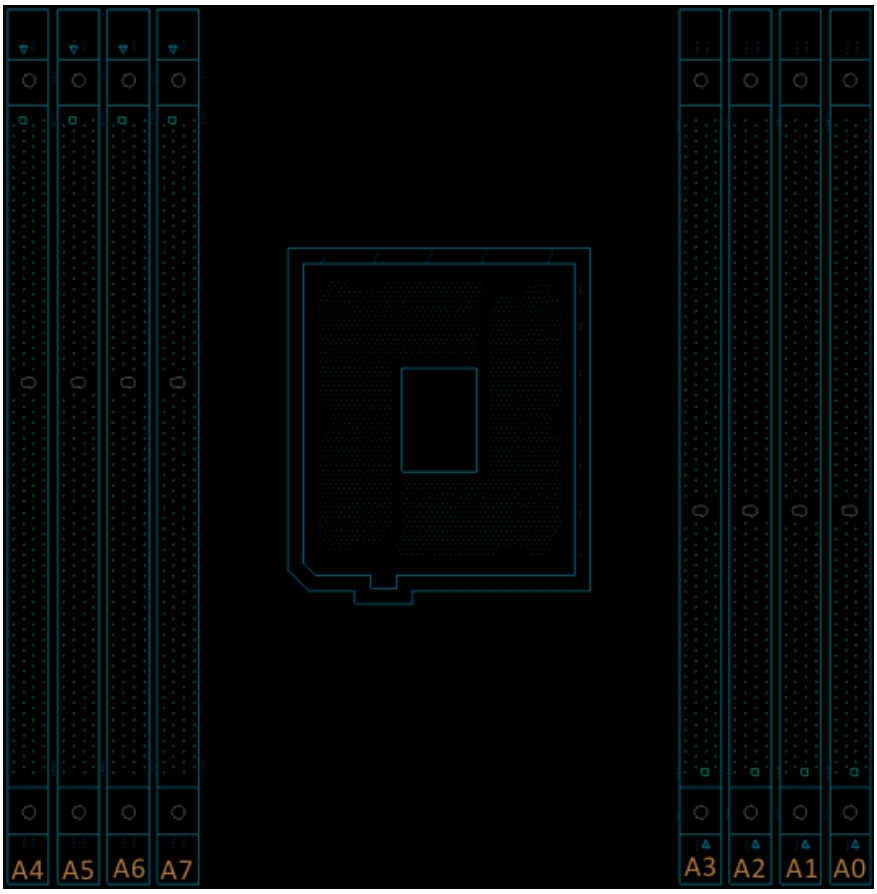

Figure 25 Efficient Performance Motherboard DIMM Numbering Silkscreen

10.8.2 Serial Console

The output stage of the system's serial console is contained on the debug card. The TX and RX signals from the SIO are sent to the debug header at the chip's logic levels (+3.3V). The debug card contains the RS-232 level shifter and the RS-232 D-9 connector.

By default, the host does console redirection through serial over LAN (SOL, see section 7.3). When the debug card is connected, debug card pin 12 is used to select console redirection between SOL and the local serial port on the card, as described in Figure 23.

10.9 Switches and LEDs

The motherboard includes a power switch, reset switch, power LED, HDD activity LED, and beep error LED.

10.9.1 Switches

The front edge of the PCB has right angle pushbutton switches. One switch is used as the system's power button the second switch is used at the system's reset button.

Note: If the ODM chooses to use smaller tactile switches, the push button actuator must be a minimum 2.5mm diameter and protrude at least 1.5mm from the switch's enclosure.

If the power switch is depressed for less than four seconds, a power management event is issued, indicating that the power switch has been triggered. If the power switch is depressed for more than four seconds, the motherboard performs a hard power off.

If the reset switch is depressed for any length of time, the motherboard performs a hard reset and begins executing the BIOS initialization code.

Each switch is identified by a label on the motherboard's silkscreen. The power button is labeled PWR and the reset button is labeled RST.

10.9.2 LEDs

The motherboard has 3 LEDs on the front edge. Figure 26 identifies each LED's color, function, and silkscreen label. The label describes the functionality of each LED.

LED Color	Function	Silkscreen Label
Blue	Power LED. This LED has the same functionality of a traditional PC power LED. It illuminates only if the motherboard is in the powered on state.	PWR
Green	Hard drive activity. This LED illuminates when there is activity on the motherboard's SATA hard drive interfaces.	HDD
Yellow	This LED replaces the functionality of the PC speaker. The motherboard causes the LED to illuminate for the same duration and sequence as the PC speaker would normally beep. The LED allows for easier diagnosis in a noisy data center environment.	BEEP

Figure 26 LED Functionality

The beep error LED patterns are described in Figure 27.

Error Description	LED Patterns						
Memory refresh timer error	On (2s)	Off (0.25s)	On (2s)	Off (0.25s)	On (2s)	Off (3s)	…(repeat)
Base memory read/write test error	On (2s)	Off (0.25s)	On (2s)	Off (0.25s)	On (0.25s)	Off (3s)	…(repeat)
Keyboard controller BAT test error	On (0.25s)	Off (0.25s)	On (0.25s)	Off (0.25s)	On (2s)		
General exception error	On (2s)	Off (0.25s)	On (0.25s)	Off (0.25s)	On (0.25s)	Off (3s)	…(repeat)
Display memory error	On (0.25s)	Off (0.25s)	On (0.25s)	Off (0.25s)	On (0.25s)		

Figure 27 Beep Error LED Patterns

11 Mechanical

Figure 28 shows the basic view of the Open Compute Project server chassis. Refer to mechanical step file provided for detailed information.

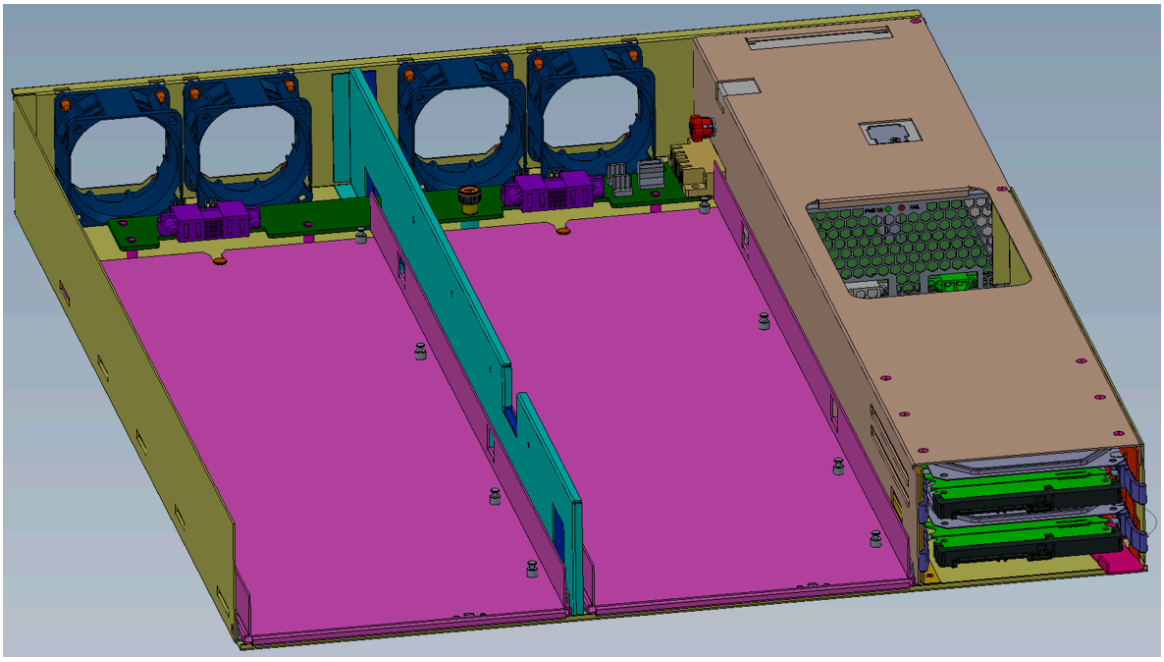

Figure 28 Open Compute Project Server Chassis for Intel Motherboards

11.1 Fixed Locations

Refer to the mechanical DXF file for fixed locations of the mounting hole, PCIe* x16 slot, and power connector.

11.2 PCB Thickness

To ensure proper alignment of the FCI power connector and mounting within the mechanical enclosure, the boards should follow the PCB stackups described in section 4.6 respectively and have 85mil (2.16mm) thickness. The mid-plane PCB thickness is also 85mil (2.16mm). The mezzanine card PCB thickness is 62mil (≈1.6mm).

11.3 Heat Sinks

The motherboard supports customized CPU heat sinks that are mounted according to the Intel specifications. The mounting device employs a backplate and receptacles for screw-down type heat sinks. The ODM must provide all keep out zones defined by Intel to ensure the heat sinks mount correctly on the board.

11.4 Silkscreen

The silkscreen is white in color and includes labels for these components:

- cpu0/cpu1
- eth0/eth1
- DIMM slot numbering, as described in 10.8.1
- LEDs, as defined in 10.9.2
- Switches, as defined in 10.9.1

11.5 DIMM Connector Color

Colored DIMM connectors indicate the first DIMM of each memory channel, whereas the remaining DIMM connectors on the same memory channel are a different color. The first DIMM on each channel is defined as the DIMM placed physically furthest from its associated CPU. This DIMM connector must be populated first when memory is only partially populated.

12 Environmental Requirements

The motherboard meets the following environmental requirements:

- Gaseous Contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Ambient operating temperature range: -5°C to +45°C
- Operating and storage relative humidity: 10% to 90% (non-condensing)
- Storage temperature range: -40°C to +70°C
- Transportation temperature range: -55°C to +85°C (short-term storage)

The full OCP system also meets these requirements. In addition, the full system has an operating altitude with no de-ratings of 1000m (3300 feet).

12.1 Vibration and Shock

The motherboard meets shock and vibration requirements according to the following IEC specifications: IEC78-2-(*) and IEC721-3-(*) Standard & Levels. The testing requirements are listed in Figure 29.

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

Figure 29 Vibration and Shock Requirements

13 Prescribed Materials

13.1 Disallowed Components

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

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

13.2 Capacitors and Inductors

The following limitations apply to the use of capacitors:

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

Only SMT inductors may be used. The use of through hole inductors is disallowed.

13.3 Component De-rating

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