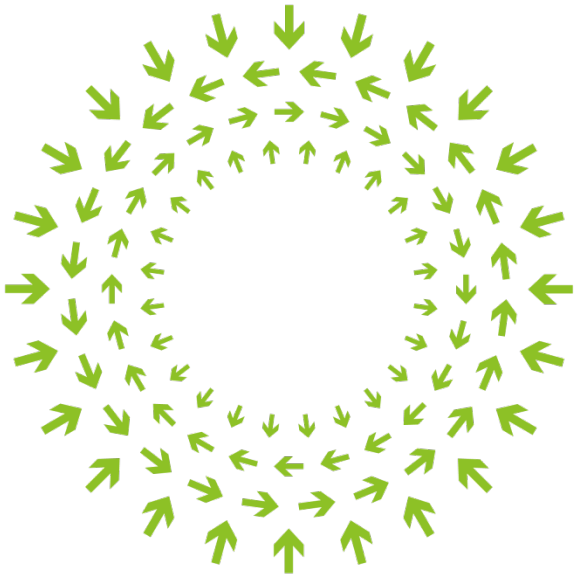

OPEN

Compute Project

High Availability Server Hardware v0.5 MB-draco-batentaban-0.5

Author: Jia Ning, Hardware Engineer, Facebook

1 Scope

This document describes the technical specifications used in the design of a server for the Open Compute Project.

2 Contents

1 Scope 2

2 Contents 2

3 Overview 3

3.1 License 3

4 System Overview 3

4.1 Common Parts with OCP Intel v2.0 Servers 3

4.2 PSU Tray 4

4.3 Midplane 4

4.4 System Assembly 6

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 high availability server with redundant power supplies.

3.1 License

4 System Overview

The server leverages the OCP Intel Motherboard Hardware Specification v2.0 (for detailed design information, see: http://opencompute.org/projects/intel-motherboard/). Instead of accommodating two server motherboard trays with one shared PSU, it accommodates one server motherboard tray with one PSU tray holding two PSUs. This adds redundant PSU capability into the server design of the Intel motherboard, to facilitate more use cases that require a redundant PSU.

4.1 Common Parts with OCP Intel v2.0 Servers

The high availability server reuses most parts of an OCP Intel v2.0 server, as described in the above specification. Any changes are detailed below.

4.2 PSU Tray

In order to support a 1+1 redundant PSU, the server uses an off-the-shelf PSU, instead of the OCP PSU. You can choose between Power One (SPAFCBK-09G) and Delta (DPS-450SB-1).

1+1 redundant PSUs are installed on the PSU tray, which is inserted into the right slot of the chassis, while the server motherboard tray is inserted into the left slot. A different PSU needs a different PSU tray with a dedicated PDB (power distribution board).

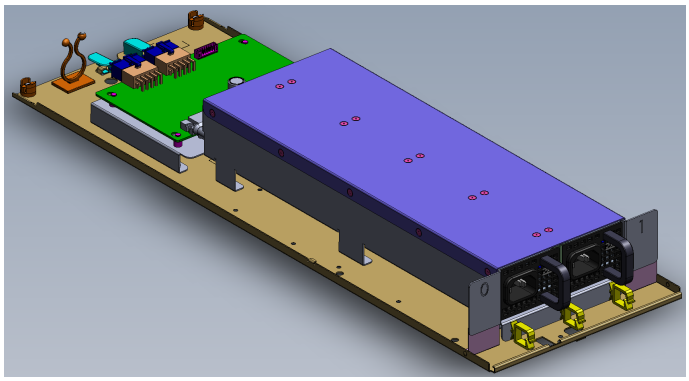

Figure 1 PSU Tray for a *Power One* PSU

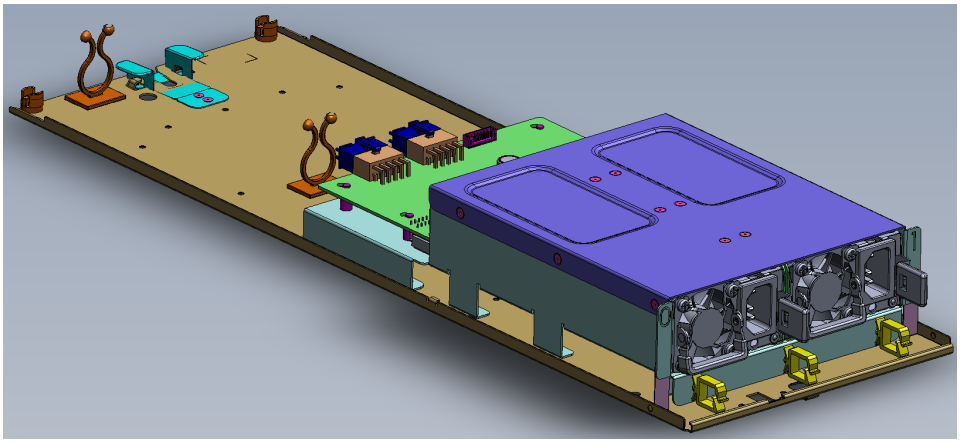

Figure 2 PSU Tray for a *Delta* PSU

4.3 Midplane

The server midplane keeps the same mechanical form factor with the same mounting locations as an OCP Intel v2.0 server system, but is modified to provide the specified interface to the PDB in the PSU tray.

4.3.1 DC Input Connection

Two power cables connect from the PDB in the PSU tray to the midplane. Each provides five wires for 12Vdc and five wires for the ground. The 12Vdc for all 10 wires comes from the same power rail, so the PDB can balance the current between both redundant PSUs. Both power cables use Molex 44476-1111 female crimp terminal or equivalent and Molex 39-01-2205 receptacle housing or equivalent, with the following pin definition. Both

cables and connectors share the same pin definition. The cables must be at least AWG #18. The 12Vdc and ground are color-coded.

Pin	Description
1~5	Ground
6~10	+12V DC

Figure 3 DC Input Cable Pin-Out

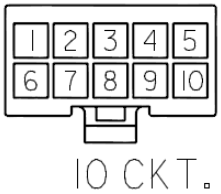

Figure 4 DC Input Cable Wiring Diagram

4.3.2 PSU Side-band Signal Connection

The side-band signal connector carries control signals and a low-speed bus between the PDB and the midplane. It has 8 pins, a 2mm pitch, with following pin definition. The connector on the midplane is Molex 89400-0820 or equivalent, mating with Molex 87369-0800 crimp housing; the cable assembly is Molex 50212 crimp terminals or equivalent. The cable is AWG #24.

The PMBUS address uses an 8-bit format, where PSU0 uses 0xB0 and PSU1 uses 0xB2. The FRU address also uses an 8-bit format, where PSU0 uses 0xA0 and PSU1 uses 0xA2.

Signal Name	Required	Direction	Description
SCLK	Yes	MB->PDB	5V tolerant, standby power, the pull-up to 5V standby is provided by the motherboard. Connects to PSU0 and PSU1 as PMBUS, to report PSU failure status. If the PSU has 3.3V standby on PMBUS, the level shifter is done on the PDB.
SDATA	Yes	<->	Same as SCLK.
Ground	Yes		Provides signal ground return between MB/midplane and the PDB.
SMBUS Alert#	Yes	PDB->MB	Low active, 5V tolerant, open drain output from the PDB. It is asserted by the PDB when any redundant PSU has a failure.
PSU_ON#	Yes	MB->PDB	Low active, open drain output from the midplane. It is asserted to turn on the PSU. It can always be asserted from the midplane so the PSU turns on by itself once AC input is available.
PSU_PWROK	Yes	PDB->MB	High active, 5V tolerant, open drain output from the PDB. It indicates the PDB's 12V output is stable.
PSU_ACOK	Optional	PDB->MB	High active, open drain output from the PDB. It indicates the PSU's AC input is OK. For the redundant power supply, a deassert in any PSU's ACOK causes a deassert of PSU_ACOK.
P5V_STBY		PDB->MB	Provide 5V standby power, 100mA.

4.3.3 Fan Connector

The midplane has four system fan connectors. The fan connector follows the *4-Wire Pulse Width Modulation (PWM) Controlled Fans Specification*, Revision 1.3, September 2005, published by Intel.

4.3.4 Motherboard Connector

The midplane has one FCI 51770-044 female right-angle power/signal connector (2P+16S+2P: 4 power blades and 16 signals), which mates with the motherboard's FCI 51730-162 male right angle header. The connector's pin definition follows the *Intel Motherboard Hardware Specification v2.0*.

4.3.5 Hot Swap Controller

One hot swap controller (ADI ADM1276) controls in-rush current and reports power consumption. The implementation follows the *Intel Motherboard Hardware Specification 2.0*.

4.4 System Assembly

The full system assembly is shown below.

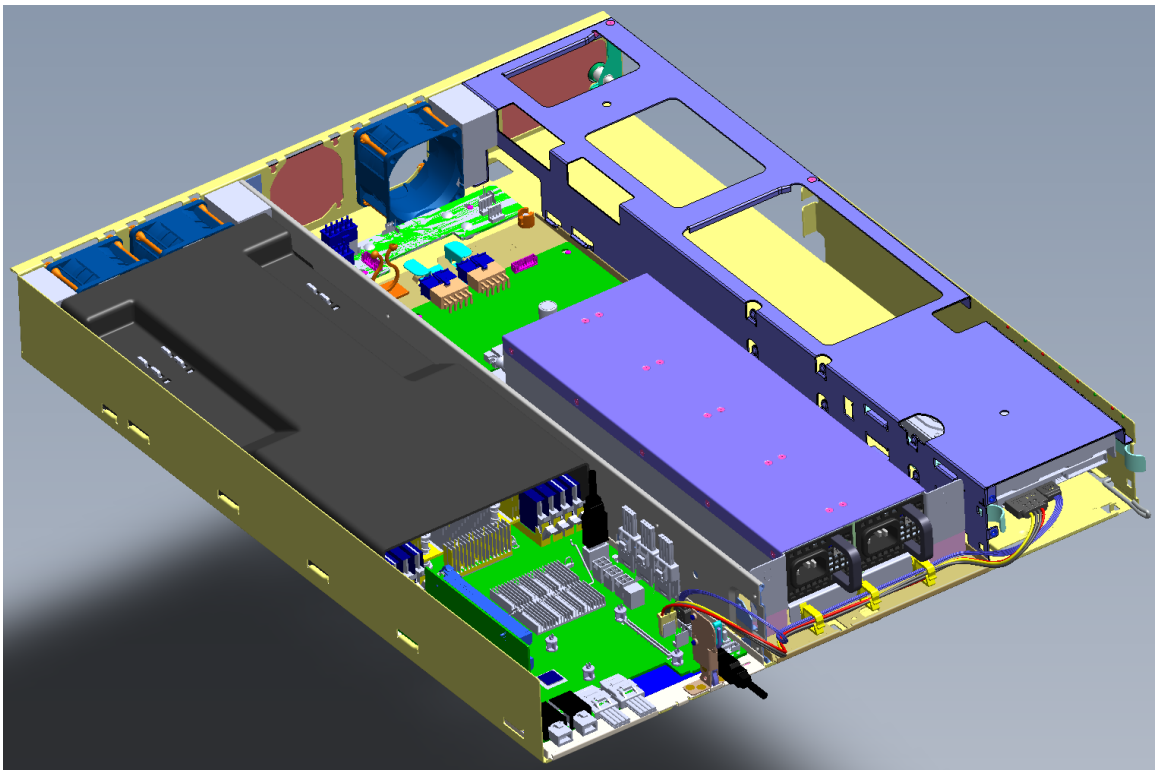

Figure 5 System Assembly