



OPEN

Compute Project

AMD Open 3.0 Modular Server

Authors: Bob Ogrey, Eugene Chung, Otto Joe

1 Revision History

Date	Revision	Description
May 2012	1.0	Initial Release to OCP
June 2012	1.1	Changed Add-in Management Controller Connector, Changed Mezzanine connector, Modified Block Diagram to support BCM5720, Added in EMC and UL safety requirements, added Power Supply Information, Integrated partner feedback
January 2013	1..0	Updated product name to AMD Open 3.0 modular server

2 Scope

This document defines the technical specifications for the general purpose AMD motherboard that will be used in financial services Open Compute Project servers.

3 Contents

1	Revision History	2
2	Scope	2
3	Contents	3
4	Overview	5
4.1	License	5
4.2	CAD Models	6
5	Product Descriptions and Requirements.....	6
6	Motherboard Features.....	7
6.1	Block Diagram for Base Design Implementation	9
6.2	Placement and Form Factor.....	9
6.3	CPU and Memory.....	10
6.4	Northbridge PCIe Usage.....	13
6.5	Southbridge/Peripheral Bus Controller	13
7	Leveraged Design	13
7.1	HPC Server in 1U Chassis	14
7.2	General Purpose Motherboard in 2U Chassis.....	16
7.3	Storage Server Motherboard in 3U Chassis	18
8	Chassis Information.....	20
8.1	Supported Configurations.....	20
8.2	Front Panel Switches and Indicators.....	20
8.3	Heat Sinks	20
9	BIOS.....	21
9.1	BIOS Chip	21
9.2	BIOS Socket	21
9.3	BIOS Source Code	21
9.4	BIOS Power Optimization Features.....	21
9.5	BIOS Setup Menu	21
9.6	Console Redirect	22
9.7	PXE Boot.....	22
9.8	Other Boot Options.....	22
9.9	Remote BIOS Update	22

9.10	Event Log	23
10	System Management	24
10.1	Requirements	24
10.2	Management Port.....	25
10.3	Temperature Sensors.....	25
10.4	Fan Connections	25
11	Power System.....	25
11.1	Power Supply	25
11.2	Power Connector Pin-out	26
11.3	Support for Open Compute Infrastructure	26
11.4	Power Budget for Open Platform Board Configurations	27
11.5	VRM Specifications	28
11.6	Power Sequencing	28
11.7	Power Optimized Hardware Features.....	28
12	I/O System.....	29
12.1	PCI-Express Slots	29
12.2	PCIe Mezzanine Card	29
12.3	IMC Management Connector	30
12.4	Network	33
12.5	USB Interfaces.....	33
12.6	SATA.....	33
12.7	Debug Header	33
13	Environmental Requirements	34
13.1	Regulatory Compliance	34
13.2	Vibration and Shock.....	34
14	Prescribed Materials	35
14.1	Disallowed Components	35
14.2	Capacitors and Inductors	35
15	Operating System Support	36
16	Adherence to AMD Motherboard Design	36

4 Overview

This document describes the "AMD Financial Services Open Platform and System Solutions," hereafter referred to as the Open Platform board.

The AMD Financial Services Open Platform is a G34 motherboard that is optimized for cost and power, and designed to fit into a 1U, 1.5U, 2U, 3U or taller chassis to support different hard drive configurations. The chassis will house this 16" by 16.7" board and come in different U-heights depending on the use case. This system is intended as a low power offering with a feature set tailored to meet the needs of large data center operators and cloud computing service providers.

The AMD Financial Services Open Platform is intended to provide a universal, highly re-useable common motherboard that targets 70% to 80% of enterprise infrastructure of Wall Street Council members. Compatibility and deploy-ability are chief design objectives and boards should fit into both the Open Compute/Facebook datacenter infrastructure as well as traditional Enterprise 19" rack enclosures. This means the solution will provide a flexible power delivery system that supports both Facebook Open Compute & Traditional Enterprise Datacenter high-efficiency power supplies. Like previous Open Compute projects, this motherboard is power-optimized and barebones, designed to provide the lowest capital and operating costs. The Open Platform board enables customer driven (versus OEM driven) hardware core features and requirements and will drive an open standards approach that should generate volume and competitive pricing.

4.1 License

As of April 7, 2011, the following persons or entities have made this Specification available under the Open Web Foundation Final Specification Agreement (OWFa 1.0), which is available at

<http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owfa-1-0>

Facebook, Inc.

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

Your use of this Specification may be subject to other third party rights. THIS SPECIFICATION IS PROVIDED "AS IS." The contributors expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the Specification. The entire risk as to implementing or otherwise using the Specification is assumed by the Specification implementer and user. IN NO EVENT WILL ANY PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CONTRIBUTORS AND LICENSORS OF THIS DOCUMENT MAY HAVE MENTIONED CERTAIN TECHNOLOGIES THAT ARE MERELY REFERENCED WITHIN THIS DOCUMENT AND NOT LICENSED UNDER THE OWF CLA OR OWFA. THE FOLLOWING IS A NON-EXHAUSTIVE LIST OF MERELY REFERENCED TECHNOLOGY: PCI Express®, Hypertransport™, and USB.

IMPLEMENTATION OF THESE TECHNOLOGIES MAY BE SUBJECT TO THEIR OWN LEGAL TERMS.

4.2 CAD Models

CAD models as fully set forth in this specification are available upon request.

5 Product Descriptions and Requirements

The following section describes the high level features for the various configurations of the Open Platform board. A single board design will be leveraged to meet the three usage cases listed below:

1. HPC server
2. General purpose server for virtualization and standard scale out
3. Storage Platform

However, in order to optimize for each usage case, a separate board sku with different component selection may be maintained. For example, the general purpose and HPC motherboard will populate only 1 of the 2 tunnel chips for IO expansion.

The common motherboard has the following features. It is a 16" x 16.7" board designed to fit into a 1U, 2U, or 3U chassis. The size of the motherboard was increased to support 3DIMMs per channel.

It has 2 processors, each with 12 memory sockets, 6 SATA connections per board, on-board 1Gb NIC with integrated management controller, up to four PCIe® expansion slots, mezzanine connector for custom module solutions, two serial port interfaces, and two USB ports. Specific PCIe card support is dependent on usage case and chassis height.

System Configuration	Description (Common motherboard design is used for each of the following system configurations)
HPC	<ul style="list-style-type: none"> • 1 DIMM per channel (U/RDDR3 1600MHz, 1866 MHz* Stretch Goal) • Fits into the 1U chassis • Cooling and Power for SE 140W parts • 1 SR5670 tunnel • Supports 6 SATA drives natively off of the "Southbridge" • Supports up to ten 2.5" total drives with add-in card (Full details in section 7.2) • Supports up to 2 low-profile PCIe cards or 1 standard height card • Single 1Gb on-board controller via BCM5725 • 10Gb solution via add-in mezzanine card
General Purpose	<ul style="list-style-type: none"> • 3 DIMM per channel (Up to 1600 MHz support) • Fits into 2U chassis • Cooling and Power for SE 140W Parts

	<ul style="list-style-type: none"> • Support for twenty five 2.5" SATA/SAS drives • Supports up to 1 standard-height and 1 low-profile PCIe cards (2 cards total) • Single 1Gb on-board controller via BCM5725 • 10Gb solution via add-in mezzanine card
Storage	<ul style="list-style-type: none"> • 3 DIMM per channel (Up to 1600 MHz support) • Fits into 3U chassis • Cooling and Power for SE 140W Parts • Support for thirty five 2.5" SATA/SAS drives • Supports up to 4 full-height, short PCIe cards • 10Gb solution via add-in mezzanine card

6 Motherboard Features

The base motherboard design supports the following feature set and will have different stuff options and implementations depending on the use case. A general description of the base motherboard design is documented below.

Component	Function
Socket	Socket G34
Processor	<ul style="list-style-type: none"> • Two sockets per board • Support for AMD processor codenamed "Magny-Cours", "Interlagos", and "Abu Dhabi" processors • Abu Dhabi ("Orochi"-Rev C) support is mandatory • Supports Infrastructure Group A, B, C: 85W, 115W, and 140W TDPs • Magny-Cours: 8/12 cores codenamed "Greyhound" for Hydra die (MCM) • Interlagos: 12/16 cores codenamed "Bulldozer" for Orochi die (MCM) • Abu Dhabi: 4/8/12/16 cores codenamed "Piledriver" (MCM) • Coherent Links: Triple x16 HyperTransport3 link supporting speeds up to 6.4 GT/s with support for HT1 operation @ 2.0 GT/s
Memory	<ul style="list-style-type: none"> • HPC config: 1 DIMM's per channel per processor (total of 4 DIMMs per processor and 8 DIMMs per system) • Support up to 1866 MHz for 1 DIMM per channel • General Purpose config: 3 DIMM's per channel per processor (total of 12 DIMMs per processor and 24 DIMMs per system) • Three RDIMMs/channel up to 1600 MT/s • 1DIMM/channel = 1600 MHz support • 2DIMM/channel = 1333 MHz support • 3DIMM/channel = 1066 MHz support • Supports DDR3 1.5V, DDR3L (1.35V), DDR3U (1.25 V) • Targeted configuration is 96 GB, 1.35V, RDDR3, 1333 MHz • Future LRDIMM (Load Reduced DIMM) support with AMD

	<p>processor codenamed "Orochi" required</p> <ul style="list-style-type: none"> UDIMM support
I/O	<ul style="list-style-type: none"> Two Tunnel SR56X0 Chipsets Single SR5670 for HPC and General Purpose Server One SR5670 and one SR5690 for Storage server Interconnect via ncHT3 16x16 link to North-Bridge. ncHT1 operation will not be supported with Orochi based products. MD SP5100 South Bridge onboard SATA 2.0 ports. 2 serial port 2 USB port interfaces. 1 via the debug header, 1 via the DP9 connector.
Drive Bays	Support for 2.5" and 3.5" drives. Number of drives dependent on configuration, use case, and ODM implementation.
Video	No embedded support.
I/O Slots	<ul style="list-style-type: none"> Up to 4 PCI Express® slots Motherboard can support up to two x8 PCIe slots and two x16 PCIe slots (4 total) Three x8 PCIe slot support in a 1U or 2U chassis
Networking	<ul style="list-style-type: none"> Embedded BCM5725. 10GB via add-in mezzanine. Population Option for BCM 5720. (No-popped once the BCM 5725 reaches production)
System Management	<ul style="list-style-type: none"> Integrated Management Controller via BCM5725. KVM support via add-in management card
System fans	Non-redundant and non-hot-swap, support cooling for Infrastructure Group A processors.
Board Dimensions	16" x 16.7"
Expected Layer Count	12 – 14 layers
Power target per MB	See Chapter 10.

6.1 Block Diagram for Base Design Implementation

Figure 1 illustrates the functional block diagram for the storage node version of the Open Platform board.

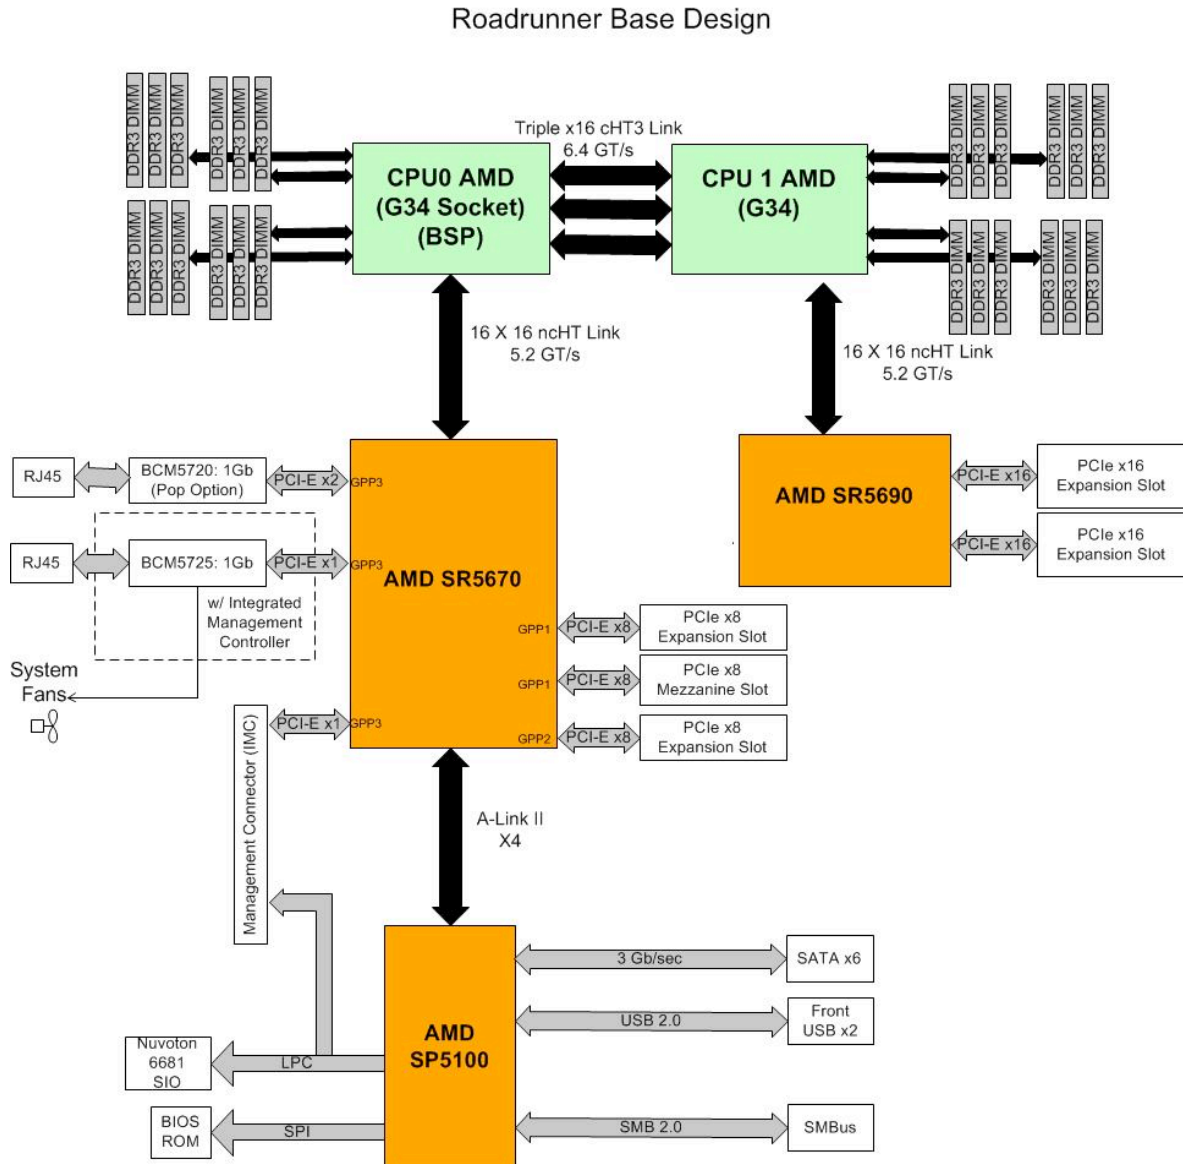


Figure 1 Functional Block Diagram

6.2 Placement and Form Factor

The motherboard's form factor is 16 x 16.7 inches with cutouts to support dual power supply locations. Figure 2 illustrates board placement. The placement shows the relative positions of key components, while exact dimension and position information will be determined in the future. Once released, the ODM should strictly follow the form factor, PCIe slot position, front IO port positions, PCIe mezzanine card connector position, power connector, and mounting holes, while other components can be shifted based on

layout routing as long as relative position is maintained. Each specific configuration (HPC, general server, or storage) will use a slightly different flavor of the Open Platform board and chassis as shown in section 7.2.




Figure 2 AMD Motherboard Component Placement

6.3 CPU and Memory

The Socket G34 processor supports multi-chip modules (MCMs). The MCMs use HyperTransport™ technology as the inter-processor link. Figure 1 illustrates the internal HT link connections between the dies and the package.

Magny-Cours (6100 Series) uses the Greyhound core while Interlagos (6200 Series) uses the Bulldozer core. Abu-Dhabi is based on the Piledriver core will be available Q2 2012 and offers drop-in compatible part with 200 MHz performance uplift.

The Open Platform board should be designed to accommodate all three G34 power bands—85W, 115W, and 140 Interlagos and Abu Dhabi parts. The SE Abu Dhabi will consume up to 125A on the core and 25A on the north bridge. The board must support a 165A IDD spike and a 125A Load step on the core rail.

Interlagos in the G34 package has the following characteristic:

- 32 nm AMD Family 15h processor
- 16, 12 or 8 CPU cores - Each pair of cores comprises a "Bulldozer" module, a.k.a Compute Unit
- Each individual core has its own 16-Kbyte L1 data cache
- Each Bulldozer module has a single 64-Kbyte L1 instruction cache
- 2MB L2 cache per Bulldozer module
- Quad-channel U/RDDR-3 memory (also supports low voltage DDR3)
- Four HyperTransport™ technology links supporting speeds up to 6.4 GT/s.
- AMD-Virtualization™ (AMD-V™) technology with Rapid
- AMD Turbo CORE technology

6.3.1 DDR3

It is a requirement to support both 1.5V, 1.35 V and 1.25V memory offerings. This requires SPD detection by the system bios to detect the correct memory voltages and hardware interfaces to program the VDDIO voltage rail based on the populated memory configuration.

Future LR-DIMM support is a requirement and can be implemented with no adverse effects to standard DDR3 operation by following the chip select routing details. This involves routing the MA3_CS_L[1:0] to all DDR3 DIMM sockets in parallel.

Surface mount DIMM connectors are recommended for this design.

4 channels DDR3 registered memory interface on each CPU:

- 3 DDR3 slots per channel per processor (total of 24 DIMMs on the motherboard)
- RDIMM/LV-RDIMM (1.5V/1.35V/1.25V), LRDIMM, and UDIMM/LV-UDIMM (1.35V/1.25V)
- SR, DR, and QR DIMMs
- DDR3 speeds of 800/1066/1333/1600/1866
- Up to maximum 768GB memory with 32GB RDIMMs

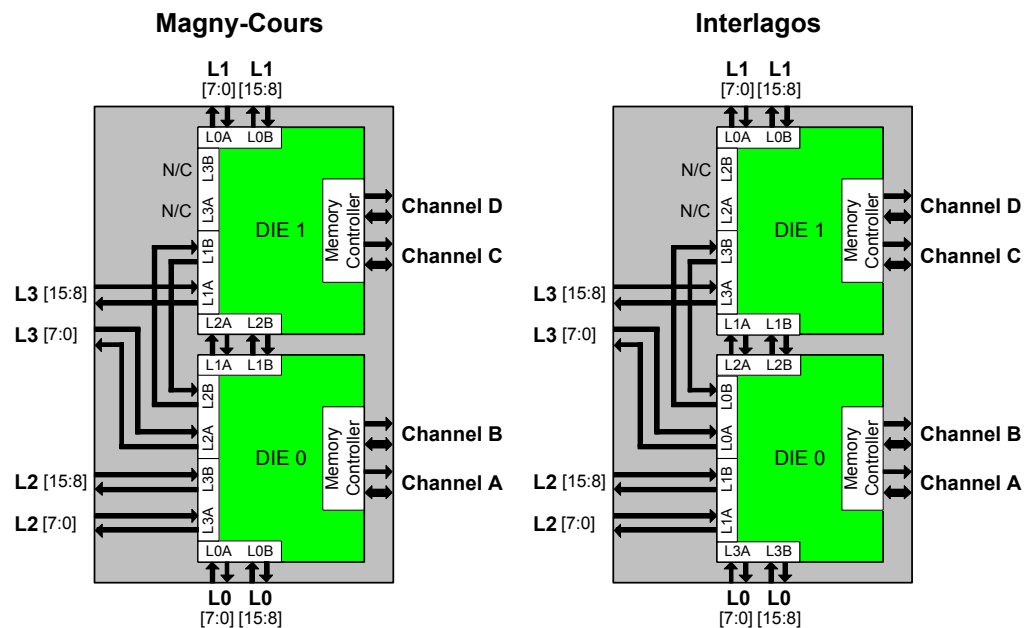


Figure 3 MCM Internal HyperTransport™ Interconnects

Figure 4 illustrates the internal memory bus connections between the dies and package.

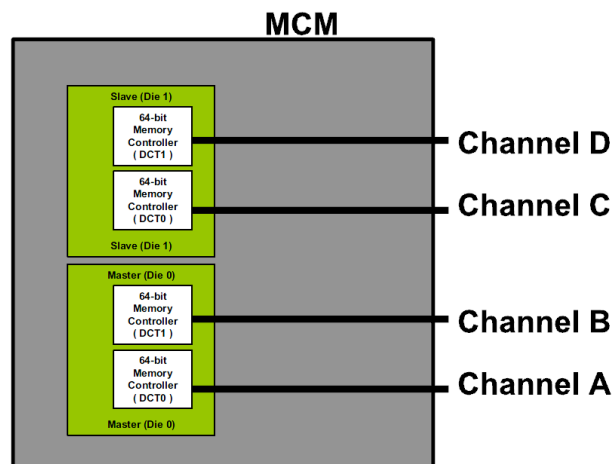


Figure 4 MCM Internal DDR Bus Connections

6.4 Northbridge PCIe Usage

The motherboard supports two SR5690 chipsets to provide maximum I/O in a balanced configuration. There are three different versions of the tunnel chipset, each with a different PCI-Express port configuration that is documented below.

	GPP1 Port	GPP2 Port	GPP3 Port	Total Available Lanes
SR5650	16 lanes	0 lanes	6 lanes	22 lanes
SR5670	16 lanes	8 lanes	6 lanes	30 lanes
SR5690	16 lanes	16 lanes	6 lanes	42 lanes

Depending on the usage case, the board will have different chipset population options. Please refer to section 6.0 for further details. For the storage configuration, a 5650 and a 5690 will be populated on the motherboard.

6.5 Southbridge/Peripheral Bus Controller

The motherboard uses the AMD SP5100 Southbridge chipset, which supports the following features:

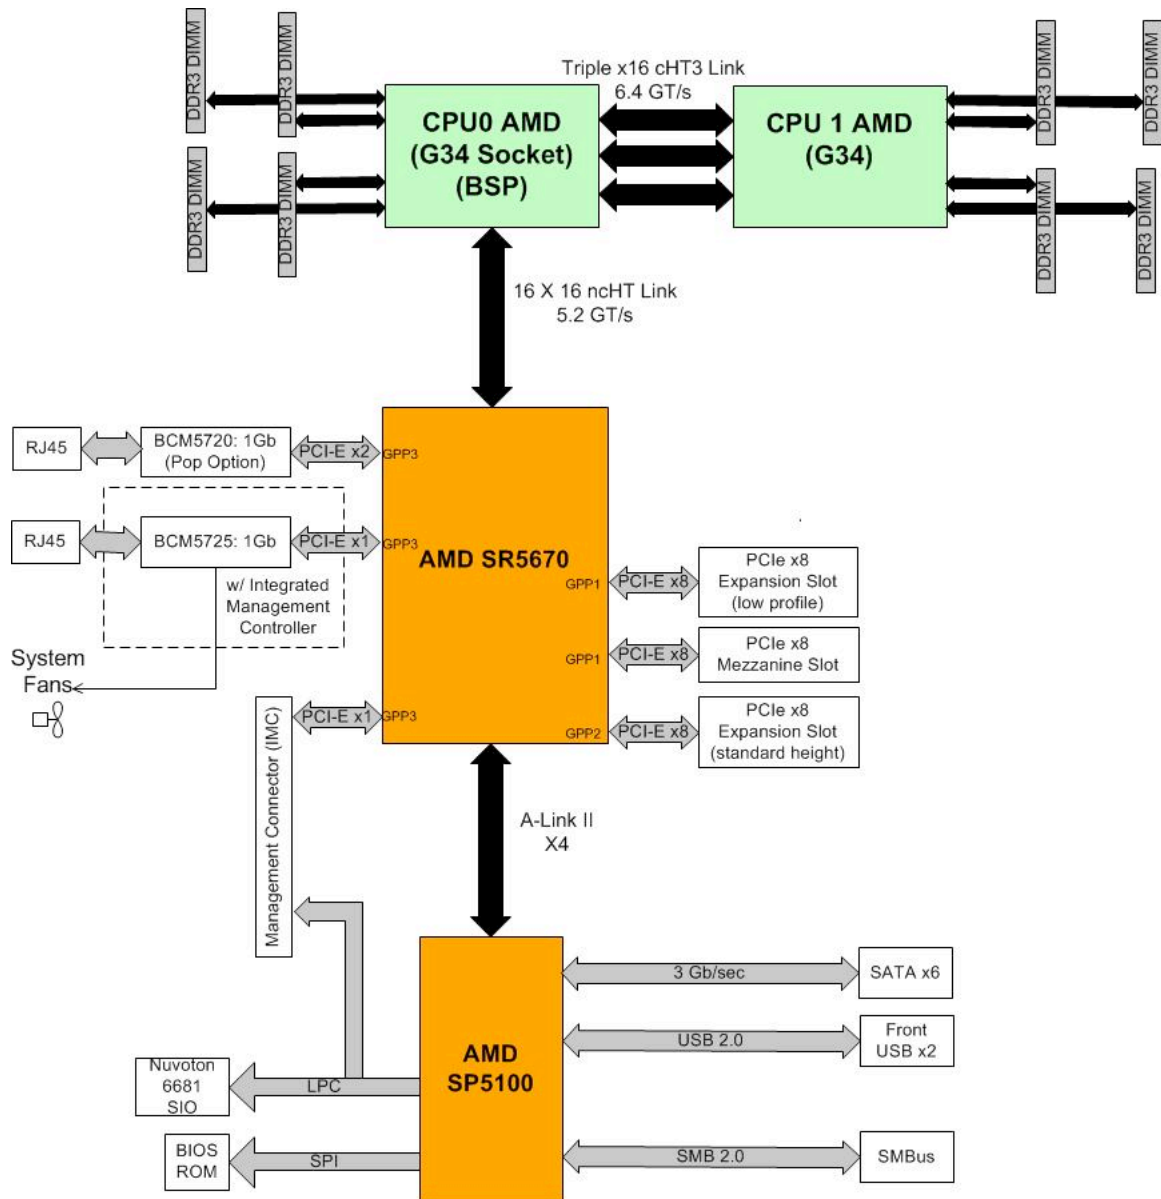
- 2 USB 2.0 ports (on the front panel)
- SATAII ports
- SPI interface
- SMBUS interface (master and slave)

7 Leveraged Design

The motherboard has been designed to be easily leveraged to support the HPC, general purpose, and storage configurations. All of the three use cases will be built around the same core design, but a few changes in both component population and the PCB will be made to optimize the design for the given application. The Base Model PCB will support two tunnel chips and 3DIMM per channel. Changes for the other configurations are listed below.

Segment	PCB	Component Change
HPC	Some memory traces are removed from the base model to support fastest possible 1DIMM per channel implementation.	Second I/O tunnel (56X0) is not populated.
General Purpose	Base Model PCB	Second I/O tunnel (56X0) is not populated.
Storage	Base Model PCB	SR5690 is populated to support two additional x16 PCIe slots. A SR5650 (in lieu of the 5690) can be populated if x8 support is sufficient.

7.1 HPC Server in 1U Chassis






Figure 5 HPC Motherboard Concept




Figure 6 1U HPC System Configuration 1

7.2 General Purpose Motherboard in 2U Chassis

Roadrunner "General Purpose" Configuration in 2U Chassis

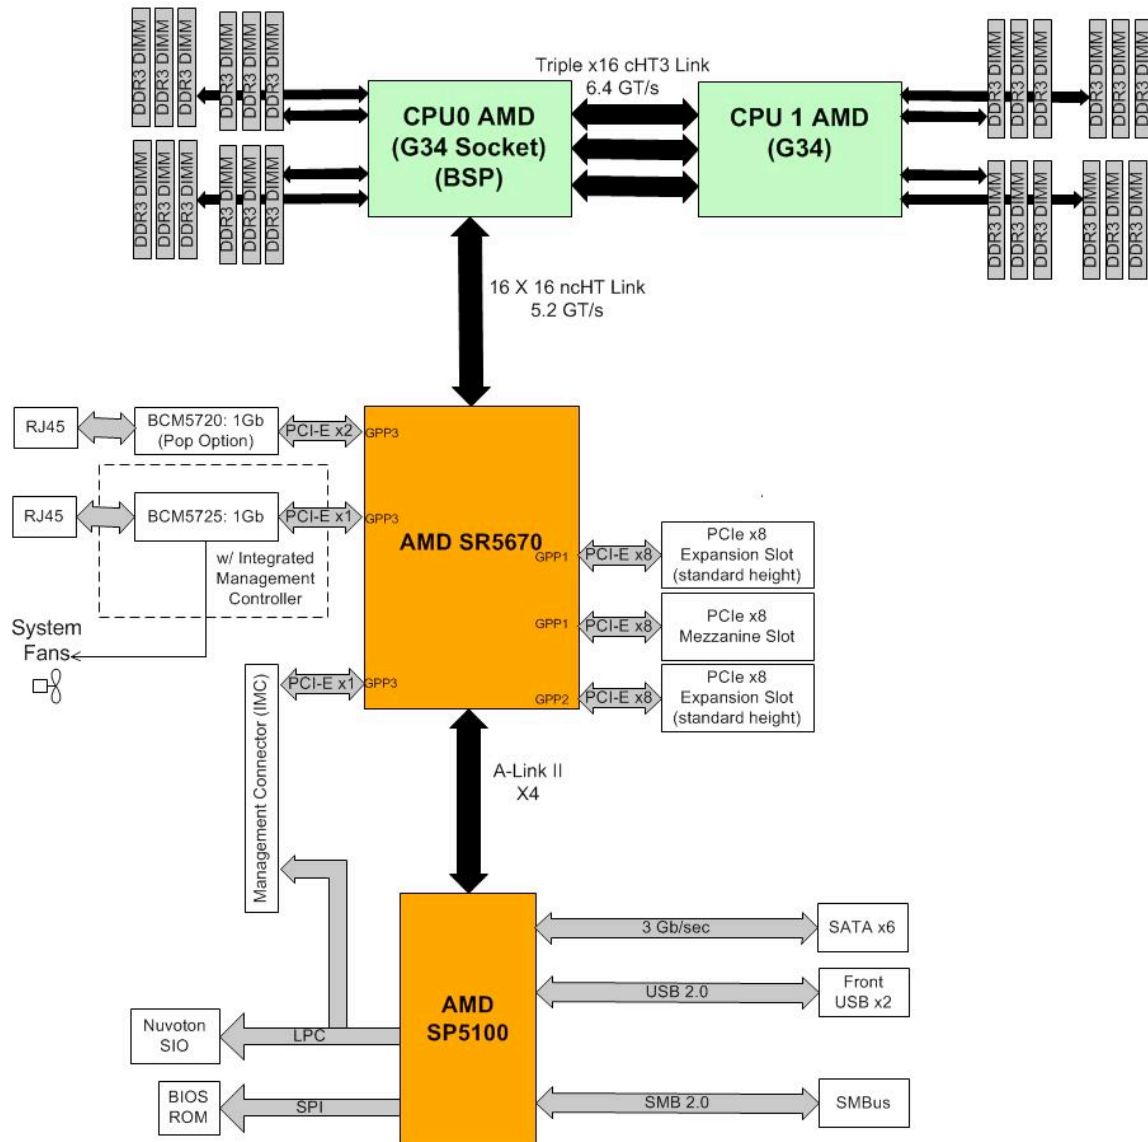


Figure 7 General Purpose Configuration in 2U Chassis




Figure 8 General Server Motherboard




Figure 9 General Server in 2U Chassis




Figure 10 Storage Server Motherboard




Figure 11 3U Storage Server System

8 Chassis Information

The Open Platform board motherboard platform is designed to fit in a custom 1U/1.5U/2U/3U chassis.

The Open Platform board can fit into an Open Rack compliant chassis that will need to be developed. It can fit into future Open Compute1.0 chassis that have been designed with multiple holes punched into the base. However, it cannot fit into currently deployed Open Compute chassis because of incompatible hole patterns.

The 1U chassis is 730mm x 438.3mm x 43.6 mm and must be compatible with standard 19" EIA rack configurations.

The chassis is designed so that it can be easily assembled and the major components of the system are readily accessible.

8.1 Supported Configurations

The chassis will be available in several different heights but will maintain the same drive cage support mounting for the various U-heights. A suggested support matrix is shown below. However, the ODM may choose to provide alternate drive cage solutions based on their existing solutions.

Financial Services Matrix				
	1U	1.5U	2U	3U
Dual G34 socket	Up to 140W SE class	Up to 140W SE class	Up to 140W SE class	Up to 140W SE class
DIMM Connectors	8 at 1DPC	24 at 3DPC*	24 at 3DPC*	24 at 3DPC*
LP PCIe	1**	3**	1**	
Standard PCIe	1**	1**	1**	4***
3.5"/2.5" HDD	4/10	4/4	12/24	12/35****
System fans	40x40x56 qty 8	60x60x38 qty 6	80x80x38 qty 8	120x120x38 qty 6
Management	Via BCM 5725 or Add in module	Via BCM 5725 or Add in module	Via BCM 5725 or Add in module	Via BCM 5725 or Add in module
Power	Redundant hot swappable, single output power supplies	Redundant hot swappable, single output power supplies	Redundant hot swappable, single output power supplies	Redundant hot swappable, single output power supplies
10GbE	Via optional Mezz card or PCIe expansion	Via optional Mezz card or PCIe expansion	Via optional Mezz card or PCIe expansion	Via optional Mezz card or PCIe expansion
SAS	Via optional PCIe expansion	Via optional PCIe expansion	Via optional PCIe expansion	Via optional PCIe expansion
SATA	6 onboard ports	6 onboard ports	6 onboard ports	6 onboard ports
** requires 1U or 1.5U riser board				
*** no riser boards required, but requires 2nd 56XX chipset to be populated				
**** preliminary HDD config estimate, actual config TBD				

8.2 Front Panel Switches and Indicators

TBD

8.3 Heat Sinks

The motherboard supports heat sinks that are mounted according to the AMD G34 heat sink specification. The mounting device employs a back-plate and receptacles for screw-down type heat sinks. The ODM must comply with all keep out zones defined by AMD.

9 BIOS

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

9.1 BIOS Chip

The BIOS uses the SP5100's SPI interface. The ODM is responsible for selecting a specific BIOS chip that meets the required functionality.

9.2 BIOS Socket

A socket on the motherboard holds the BIOS chip, which allows for manual replacement of the BIOS chip. The BIOS socket is easily accessible; other components on the motherboard do not interfere with the insertion or removal of the BIOS chip.

9.3 BIOS Source Code

The BIOS source code comes from AMI or Phoenix. The ODM is responsible for maintaining the BIOS source code to make sure it has latest code release from AMI or Phoenix and AMD.

9.4 BIOS Power Optimization Features

The BIOS is tuned to minimize system power consumption and should expose all device configuration, device feature, and power saving options in BIOS, to provide maximum flexibility to tune the system to the workload and operational environment. It should enable the following features:

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

Best practices include the following:

- Full C1E support.
- "Efficient single fan variant" system fans that are non-redundant and utilize an aggressive fan management algorithm
- Disable HT Assist (Probe Filter Option)
- BIOS F2 setting(s) for HT Link Frequency. Enable cHT1 operation by default
- BIOS F2 setting(s) for HT link width. Enable 8-bit coherent link by default.
- Bios option to disable IOMMU operation
- Bios disable option for the BCM5275 NIC ports. '
- PSI options enabled by default
- Power Save options enabled
- A "green button" bios option to load all power optimized features
- Bios should enable CC6 (per core power gating) features in the processor. No specific HW changes need to be made to support this feature.

9.5 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:

- Setting for power feature after AC failure; default is set to keep last state.
- Setting for console redirection. Selectable options to support select console redirection from local COM port or the BMC's virtual UART for SOL.
- 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.
- Display power on self test (POST) results during boot up.
- If a CMOS checksum error happens (for example, caused by a BIOS update), the BIOS loads the system default automatically after showing a text message in the console for 5 seconds and rebooting the system to apply the update without user input.
- Setting to disable all "wait for keyboard input to continue" features.

9.6 Console Redirect

The BIOS detects the presence of a video card in a PCIe slot. If a video card is present, the BIOS directs its output to the video card. If no video card is present, the BIOS directs its output to the board-mounted RS-232 console output.

9.7 PXE Boot

The BIOS supports PXE boot. When PXE booting, the system first attempts to PXE 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).

9.8 Other Boot Options

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

9.9 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 tool(s) should have the following capabilities:

- Update from the operating system over the LAN – the OS standard is CentOS v5.2
- Can complete update with a single reboot (no PXE boot, no multiple reboots)
- BIOS update or BIOS setup option change take no more than 5 minutes to complete
- No user interaction (like prompts)
- Can be scripted and propagated to multiple machines

9.10 Event Log

The BIOS logs system events through the baseboard management controller (BMC).

9.10.1 Logged Errors

- **CPU/memory errors:** Both correctable ECC and uncorrectable ECC errors should be logged into event log. Error categories include DRAM, HyperTransport Link, and L3 Cache.
- **HyperTransport 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.
- **Internal parity errors:** All errors which have status register should be logged into the event log. Fatal, non-fatal, or correctable classification follows the chipset vendor's recommendation.
- **PCIe errors:** All errors which have status register should be logged into 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 should be logged into the event log.
- **Power errors:**
- **MHOT and PROCHOT errors:** MEMHOT events should be logged with event source information indicating whether the event was triggered by a DIMM or a DIMM's Voltage Regulator. PROCHOT events should be logged with event source information indicating whether the event was triggered by a CPU or the CPU's Voltage Regulator.

9.10.2 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 Facebook debug card.
- **HyperTransport errors:** Follow the chipset vendor's suggestion.
- **PCIe errors:** Follow the chipset vendor's suggestion.

10 System Management

The following sections describe the system management features implemented on the Open Platform board platform.

10.1 Requirements

- Out-of-band Management is provided by the integrated BMC (baseboard management controller) in the BCM5725 NIC.
- A separate NCT6681D eSIO/Hardware Monitor assists the BMC to provide fan speed control and fan/temperature monitoring.
- The integrated BMC provides out-of-band remote management access and control via: 1) the GbE network port in the BCM5725 or via a standard NC-SI interface that can be connected to a separate discrete NIC.
- The integrated BMC provides an IPMI/DCMI event log. BIOS will store memory or boot errors in the event log. The BMC will store sensor threshold events in the event log. The event log can be read in-band or out-of-band using IPMI/DCMI standard interfaces.
- The BMC provides the following **out-of-band** remote control features:
 - Support for IPMI and DCMI network interfaces
 - Power-on/off and hard reset power control (via IPMI/DCMI or SMASH-2.0/WSManagement protocols)
 - Serial text console redirect (via IPMI SOL protocol, telnet, or SSH protocols)
 - IPMI Event log (via IPMI/DCMI)
 - Temperature sensors and hardware inventory (via IPMI/DCMI)
- The BMC provides the following **in-band** interfaces for systems management agents running on the host operating system or BIOS
 - KCS host HW/SW interface compliant w/ DCMI and IPMI
 - IPMI Event log rd/wr (via IPMI/DCMI commands)
 - Sensor and temperature monitoring (via IPMI/DCMI commands)
 - BMC Configuration (via IPMI/DCMI commands)
- The BMC also provides the following **advanced** management features:
 - Power capping per DCMI specification
 - Boot to a remote network CD image
 - Active Directory Authentication via Kerberos for SMASH 2.0/WSManagement and SSH console redirect out-of-band protocols
 - Web browser GUI for basic management functions

Notes:

- DCMI (Data Center Management Interface) is a standard targeted for Data Center management. It uses interfaces defined in IPMI as the underlying protocol. DCMI minimizes the number of optional interfaces, requires support for specific temperature sensors, and adds power capping control to IPMI.
- IPMI (Intelligent Platform Management Interface) is a widely used legacy management interface for servers. It uses a simple UDP-based network protocol and security protocols unique to IPMI.

- SMASH (System Management Architecture for System Hardware) is the latest DMTF-defined management protocol. SMASH 2.0 defines a web-service protocol stack that uses SOAP-formatted messages, WSManagement, HTTP/S, TCP protocol stack with well accepted transport security.

10.2 Management Port

The onboard RJ45 can be configured as a dedicated management port or a shared 1-port GbE port. An on-board NC-SI connector will be populated on Ajax to support standard sideband interface to any add-in LAN cards.

10.3 Temperature Sensors

The motherboard should support these thermal sensors:

- Two to monitor temperatures for CPU0 and CPU1, retrieved through the CPU's temperature sensor interface (TSI)
- 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

10.4 Fan Connections

The motherboard has fan tachometer and PWM connections to control the fans listed in section 7.1.

11 Power System

This chapter describes the power sub-system for the Open Platform board platform.

11.1 Power Supply

The silver box should be specified and designed to provide the highest efficiencies possible. The chassis will be designed such that it supports a redundant power supply.

In addition, the power supply should meet the following criteria:

- Power supply should have an minimum 90% efficiency
- Power supply input should be TBD
- Multi-output power supply output
- PMBus interface support
- Redundant and hot-swappable power supply support

The board outline and initial specification was modeled using the Lite-on PS PS2112 supply. The output connector is card edge extension of PCB and blind mates with 32 position each side (Tyco 1761469(vertical) and 1761468(right angle) or FCI pn 10046971-100LF (vertical) and 10053363-200LF (right angle). The mechanical design uses a straddle-mount version of these connectors (Part Number TBD).

11.2 Power Connector Pin-out

PCB Top Side	
<i>Pin #</i>	<i>Signal Name</i>
53-64	+12V
41-52	RTN
40	RS+
39	12v_SB
38	PS_A0
37	POK
36	Return
35	SCL
34	-PS_Present
33	SDA
PCB Bottom Side	
<i>Pin #</i>	<i>Signal Name</i>
12-Jan	+12V
13-24	RTN
25	NA
26	RS-
27	Vin_Good
28	Cshare
29	-PS_On
30	PS_Kill
31	Reset
32	Alert

11.3 Support for Open Compute Infrastructure

The Open Compute Infrastructure supports three power bus bars that run along both sides of the back of an Open Compute Rack.

The Open Platform board will fit seamlessly into the open compute infrastructure by utilizing a PCB adapter (red board in the pictures below) that connects to the bus bar. In addition to interfacing with the open rack bus bars, this PCB will take the place of the silver box supplies and may provide any needed intermediate power conversion, such as 3.3V generation from 12V.

A conceptual drawing is shown below of Ajax in the Open Compute high PUE power distribution system.




Figure 12 Open Compute Support for the Open Platform Board

11.4 Power Budget for Open Platform Board Configurations

The summarized power budget below is shown for three different configurations—HPC, General Purpose, and Storage. HPC has the highest performing, highest TDP CPUs, limited memory and IO. The 1U configuration is also the hardest configuration to cool and requires the highest power fans.

The general purpose configuration has more IO, standard power parts, fully populated DIMM slots, and uses SATA hard drives. The storage node is configured with 3 SAS controllers and fully populated memory.

It is important to note that these are maximum values. During production operation, bios controls can turn off features to significantly reduce power consumption from idle to 100%.

It is expected that the 3U storage node will use a different power supply than the 1U and 2U configurations.

		pwr	qty	SE Power		pwr	qty	Standard Power		pwr	qty	Standard Power	
Processors		140	2	280.0		115	2	230.0		115	2	230	
Ajax + CPU+ Memory + IO				376.8				422.8				483.5	
2.5" SSD		6	1	6.0									
SATA HDD	2.5" SATA					5.0	20	100.0		5.0	0	0	
SAS HDD	2.5" SAS 15k RPM	16.3	0	0.0		16.3	0	0.0		16.3	35	570.5	
Fans*		21.4	8	170.9		8.6	8	69.1		12	6	72	
System power (Twin + Ajax) - HPC Node				553.6	General Purpose-2U				591.9	Storage Node-3U			1,126.0

Figure 13 Power Budget for the Open Platform Board

11.5 VRM Specifications

The VRM and motherboard should support PSI features that allow the VRM to run in its peak efficiency range, i.e. redundant phases on the VRM will be turned off when the processor is operating at a low power point. The CPU VRM solutions should be load line enabled and designed to fit the specified power loads only.

In addition, the voltage regulators with the following characteristics are recommended:

- Load line to minimize load step transient.
- Differential Sense for accurate voltage sensing.
- Support for high value, low ESR ceramic caps and low ripple noise.
- In general, the target for ripple from the regulator should be < 10% of the maximum allowable voltage range. This allows 90% for all other sources of noise.

11.5.1 Target Power Efficiencies

The following are the target power efficiencies for each given group:

- Core VR: > 90% across the full operating load
- DDR VR: > 90% across the full operating load
- Silverbox: > 90% depending on load

11.6 Power Sequencing

The Open Platform board should follow power sequencing requirements for the G34, SP5650, and SP5100. Specific sequencing requirements are outlined in available design guides.

11.7 Power Optimized Hardware Features

Section 5.2 and Section 5.3 describes the high level power optimized hardware and BIOS features needed for the Open Platform board.

11.7.1 Switching Regulators versus Linear Regulators

Because the Open Platform board is designed to maximize system power savings, careful consideration must be undertaken when designing power delivery to the various system components. There are cost and board space considerations to balance, but it is generally recommended that all voltage rails that consume more than 1W of power be implemented using more efficient switching solutions.

Specific space, cost, and efficiency tradeoffs should be carefully analyzed during the earliest stages of the design.

11.7.2 VDDR Implementation

In the past, VDDR and VLDT often shared a single regulator since both rails were specified for 1.2V operation. However, for low power designs, VDDR and VLDT should implement two separate regulators.

12 I/O System


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

12.1 PCI-Express Slots

A board fully populated with both SR56X0 chipset can support two x8 slots and two x16 slots. The second tunnel chip is a stuff option for heavy I/O configurations. If populated with a 5690, the second tunnel provides enough lanes to feed two x16 slots. If populated with a 5650, two x8 slots are available. Section 7.1 describes the mechanical limitations (riser needed, low profile vs. standard, etc.) for each chassis height.

12.2 PCIe Mezzanine Card

The motherboard has one mezzanine card connector that is intended for custom and modular 10Gb solutions. The mezzanine connector pin out and board outline are shown below and leverages the pin-out and board dimensions used in Open Compute 2.0 platforms.



Mezzanine Connector 120pin x1					
Pin Name		A		Pin Name	
P12V	61	1		MEZZ_PRSENT1_N	
P12V	62	2		P5V_AUX	
P12V	63	3		P5V_AUX	
GND	64	4		P5V_AUX	
GND	65	5		GND	
P3V3_AUX	66	6		GND	
GND	67	7		P3V3_AUX	
GND	68	8		GND	
P3V3	69	9		GND	
P3V3	70	10		P3V3	
P3V3	71	11		P3V3	
P3V3	72	12		P3V3	
GND	73	13		P3V3	
SMB_LAN_3V3STB_ALERT_N	74	14		RSVD (MEZZ_CPRSNT1_N)	
SMB_LAN_3V3STB_CLK	75	15		RSVD (MEZZ_CPRSNT2_N)	
SMB_LAN_3V3STB_DAT	76	16		RSVD (SSD_PRSENT_N)	
PCIE_WAKE_N	77	17		RST_PLT_MEZZ_N	
RSVD(DA_DSS)	78	18		RSVD (MEZZ_SMCCLK)	
GND	79	19		RSVD (MEZZ_SMDATA)	
RSVD (SATA_TX+)	80	20		GND	
RSVD (SATA_TX-)	81	21		GND	
GND	82	22		RSVD (SATA_RX+)	
GND	83	23		RSVD (SATA_RX-)	
CLK_100M_MEZZ2_DP	84	24		GND	
CLK_100M_MEZZ2_DN	85	25		GND	
GND	86	26		RSVD (CLK_100M_MEZZ1_DP)	
GND	87	27		RSVD (CLK_100M_MEZZ1_DN)	
MEZZ_TX_DP_C<0>	88	28		GND	
MEZZ_TX_DN_C<0>	89	29		GND	
GND	90	30		MEZZ_RX_DP<0>	
GND	91	31		MEZZ_RX_DN<0>	
MEZZ_TX_DP_C<1>	92	32		GND	
MEZZ_TX_DN_C<1>	93	33		GND	
GND	94	34		MEZZ_RX_DP<1>	
GND	95	35		MEZZ_RX_DN<1>	
MEZZ_TX_DP_C<2>	96	36		GND	
MEZZ_TX_DN_C<2>	97	37		GND	
GND	98	38		MEZZ_RX_DP<2>	
GND	99	39		MEZZ_RX_DN<2>	
MEZZ_TX_DP_C<3>	100	40		GND	
MEZZ_TX_DN_C<3>	101	41		GND	
GND	102	42		MEZZ_RX_DP<3>	
GND	103	43		MEZZ_RX_DN<3>	
MEZZ_TX_DP_C<4>	104	44		GND	
MEZZ_TX_DN_C<4>	105	45		GND	
GND	106	46		MEZZ_RX_DP<4>	
GND	107	47		MEZZ_RX_DN<4>	
MEZZ_TX_DP_C<5>	108	48		GND	
MEZZ_TX_DN_C<5>	109	49		GND	
GND	110	50		MEZZ_RX_DP<5>	
GND	111	51		MEZZ_RX_DN<5>	
MEZZ_TX_DP_C<6>	112	52		GND	
MEZZ_TX_DN_C<6>	113	53		GND	
GND	114	54		MEZZ_RX_DP<6>	
GND	115	55		MEZZ_RX_DN<6>	
MEZZ_TX_DP_C<7>	116	56		GND	
MEZZ_TX_DN_C<7>	117	57		GND	
GND	118	58		MEZZ_RX_DP<7>	
GND	119	59		MEZZ_RX_DN<7>	
MEZZ_PRSENT2_N	120	60		GND	

Expected solutions to be evaluated include the Broadcom 57810, Mellanox Connect-X-2, and Solar Flare SFC9000. The Mellanox solution is an x8 Gen2 PCIe solution, while the Broadcom 57810 is a Gen3 PCIe solution. Solar Flare is an x8, Gen2 PCIe solution.

12.3 IMC Management Connector


The server provides support for an add-in management card solution that plugs into right angle connector on the motherboard. The motherboard connector is a 230-pin,

0.5mm pitch, right angle connector from Foxconn with a stack height of 7.8mm. The manufacturer part number is ASOB326-S78N-7F.



This connector accepts off the shelf management cards that adhere to the following pin-out and board outline. These management cards provide KVM support or basic server management in lieu of the BCM5725.

12.3.1 Board Outline



12.3.2 Pin-out

The pin-out is shown below. Additional signal descriptions are detailed separately in the AMD Management Connector Interface Design Guide.

Signal Name	Pin #	Pin #	Net Name
5V_RUN	1	2	5V3_RUN
5V_RUN	3	4	5V3_RUN
5V3_DUAL	5	6	5V3_RUN
5V3_DUAL	7	8	5V3_RUN
GND	9	10	GND
GND	11	12	GND
GND	13	14	GND
GND	15	16	GND
5V3_DUAL	17	18	5V_DUAL
5V3_DUAL	19	20	5V_DUAL
5V3_DUAL	21	22	5V_DUAL
5V3_DUAL	23	24	5V_DUAL
KEY	25	26	KEY
KEY	27	28	KEY
KEY	29	30	KEY
KEY	31	32	KEY
KEY	33	34	KEY
KEY	35	36	KEY
PCIE_RXN	37	38	PCIE_WAKE_L
PCIE_RXP	39	40	PCIE_TXN
GND	41	42	PCIE_TXP
PCIE_REFCLKN	43	44	GND
PCIE_REFCLKP	45	46	PCIE_RST_L
GND	47	48	PCIE_PRSTN1_L
RSVD	49	50	PCIE_PRSTN2_L
LPC_FRAME_L	51	52	RSVD
LPC_LAD0	53	54	GND
LPC_LAD1	55	56	RSVD
GND	57	58	LPC_PME_L
LPC_LAD2	59	60	LPC_SERIRQ
LPC_LAD3	61	62	LPC_CLK
LPC_RST_L	63	64	GND
LPC_SMI_L	65	66	RSVD
RSVD	67	68	RSVD
RSVD	69	70	RSVD
GND	71	72	RSVD
RSVD	73	74	RSVD
RSVD	75	76	RSVD
GND	77	78	GND
RSVD	79	80	RSVD
RSVD	81	82	RSVD
GND	83	84	RSVD
RSVD	85	86	RSVD
RSVD	87	88	GND
GND	89	90	RSVD
RSVD	91	92	RSVD
RSVD	93	94	RSVD
GND	95	96	RSVD
RSVD	97	98	GND
RSVD	99	100	RSVD
GND	101	102	RSVD
RSVD	103	104	RSVD
RSVD	105	106	RSVD
GND	107	108	GND
SMBus_SCL1	109	110	SMBus_SCL2
SMBus_SDA1	111	112	SMBus_SDA2
SMBus_ALERT1_L	113	114	SMBus_ALERT2_L
GND	115	116	GND
RSVD	117	118	RSVD
RSVD	119	120	RSVD
GND	121	122	RSVD
USB0_P	123	124	GND
USB0_N	125	126	USB1_P
GND	127	128	USB1_N
USB_CLK	129	130	GND
GND	131	132	RSVD
RSVD	133	134	RSVD
RSVD	135	136	RSVD
RSVD	137	138	RSVD
RSVD	139	140	RSVD
RSVD	141	142	RSVD
RSVD	143	144	RSVD
RGMIU_TX_CLK	145	146	RSVD
GND	147	148	GND
RMII_TXD0	149	150	RSVD
RMII_TXD1	151	152	RSVD
RGMIU_TXD2	153	154	RSVD
RGMIU_TXD3	155	156	RSVD
RMII_TX_EN	157	158	GND
GND	159	160	CARD_ID0
RMII_RX_DV	161	162	CARD_ID1
RMII_RXD0	163	164	CARD_ID2
RMII_RXD1	165	166	IMC_CARD_PRE5_L
RGMIU_RXD2	167	168	RSVD
RMII_RX_ER	169	170	RMII_REF_CLK
GND	171	172	GND
RGMIU_RX_CLK	173	174	RMII_MODE
RGMIU_MODE	175	176	RMII_MDC
RSVD	177	178	RSVD
RSVD	179	180	RSVD
RSVD	181	182	RSVD
RSVD	183	184	RSVD
RSVD	185	186	RSVD
RSVD	187	188	RSVD
GND	189	190	GND
RSVD	191	192	RSVD
RSVD	193	194	RSVD
GND	195	196	RSVD
RSVD	197	198	RSVD
SYS_ACPI_STATE0	199	200	GND
SYS_ACPI_STATE1	201	202	RSVD
SYS_ACPI_STATE2	203	204	RSVD
VGA_VSYNC	205	206	RSVD
VGA_VSYNC	207	208	SYS_PWRBTN_L
VGA_RED	209	210	SYS_RSTBTN_L
GND	211	212	SYS_NMI0BTN_L
VGA_GREEN	213	214	SYS_SMI_L
GND	215	216	SYS_ALL_PWROK
VGA_BLUE	217	218	GND
GND	219	220	SYS_PCI_RST_L
VGA_BDC_CLK	221	222	SYS_IO_EXP_INT_L
VGA_BDC_DATA	223	224	MCARD_PWRBTN_L
NCSI_TXD0	225	226	MCARD_RSTBTN_L
NCSI_TXD1	227	228	MCARD_NMI0BTN_L
NCSI_TX_EN	229	230	MCARD_LOCAL_LOCK_L
NCSI_RX_DV	231	232	RSVD
GND	233	234	MCARD_CLR_CMOS_L
NCSI_RXD0	235	236	MCARD_SCL_INT_L
NCSI_RXD1	237	238	GND
NCSI_RX_ER	239	240	FAULT_LED_L
NCSI_REFCLK	241	242	CHASSIS_ID_L

12.4 Network

The motherboard has one BCM 5275 LAN chip on board to support the RJ45 connector. The BIOS supports PXE boot on the RJ45 port.

There is also a stuff option to provide support for the BCM5720. The BCM5720 is x2 PCI-Express 2.0 controller and has its own RJ45 connector on the motherboard. The NC-SI sideband interface connects to the management controller connector and is used to provide a management port if the BCM 5275 is not populated.

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.

12.5 USB Interfaces

The motherboard has two external USB ports located and the BIOS should support the following USB devices:

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

12.6 SATA

The motherboard has SP5100 interfaces on board, which support up to six SATA ports.

12.7 Debug Header

The motherboard includes a debug header on the front of the motherboard to display POST codes. 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. The debug card has a key to match with the notch to avoid pin shift when plugging in.




Figure 14 Debug Header

Pin (CKT)	Function
1	Low HEX character [0] LSB
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=SQL; 0=local)
13	GND
14	VCC (+5VDC)

Figure 15 Debug Header Pin Definition

13 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°.
- 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).

13.1 Regulatory Compliance

The 1U, 2U, and 3U server platform implementations must also meet the following regulatory compliances:

- FCC/CE Class "A" EMC
- UL Safety Enabled

13.2 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 below.

	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 16 Vibration and Shock Requirements

14 Prescribed Materials

14.1 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

14.2 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 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 motherboard).

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

15 Operating System Support

Most up to date information about OS support is listed below.

ASSUMES latest updates/patches are installed*		
Enabled Optimized to support some or all of "Bulldozer's" new features	Compatible Will boot and run but not take advantage of "Bulldozer's" new features outside of new instructions	Not Supported Will not run on "Bulldozer" platforms and/or will not be supported by OSV
<p>Includes new instruction support:</p> <ul style="list-style-type: none"> Linux kernel 2.6.37 + , 3.0 + Novell SLES 11 SP2 (includes Xen) RHEL 6.2 with KVM (with latest z-stream patches) Windows Server 2008 R2 SP1 (optional scheduler patch available) Windows Server 2012/Hyper-V (in development) Xen 4.1 + Ubuntu 11.04 (w/ KVM) VMware vSphere 5.0 <p>Versions in this category also include latest software advances</p>	<p>Includes new instruction support:</p> <ul style="list-style-type: none"> Linux kernel 2.6.32 – 2.6.36 Novell SLES 11 SP1 RHEL 6.1 Ubuntu 10.10 <p>Does <u>not</u> support new instructions for either Bulldozer or Sandy Bridge:</p> <ul style="list-style-type: none"> Hyper-V R1 Hyper-V R2, Hyper-V R2 SP1 Novell SLES 10 SP4 and higher RHEL 5.7 (included KVM) Solaris 10u9, 11 VMware vSphere 4.1u2 Windows Server 2003 R2 SP2 Windows Server 2008 R2 Windows Server 2008 SP2 Xen 3.4.2 <p><i>Will run but not necessarily provide performance uplift</i></p>	<ul style="list-style-type: none"> Linux kernel 2.6.31 or earlier Novell SLES 10 thru SP3 Novell SLES 11 RHEL 4.x RHEL 5.0 – 5.5 RHEL 5.6 (can run with patches but is not supported by Red Hat) RHEL 6.0 Solaris 10 – 10u8 VMware ESX 3.5 VMware ESX 4.0 – 4.1u1 Windows Server 2003 versions prior to R2 SP2

Figure 17 Supported Operating Systems

16 Adherence to AMD Motherboard Design

The expectation is that close partnership between AMD and a selected ODM will occur. This will include joint review of specification, schematic, stack-up, layout, thermal, and chassis design.