

OPEN
Compute Project

OCP U.S. SUMMIT 2017

Santa Clara, CA



Facebook OCP 2S Server Tioga Pass

Whitney Zhao, Hardware Engineer, Facebook

OPEN HARDWARE. OPEN SOFTWARE. OPEN FUTURE.



Agenda

1

Overview

2

What's New

3


Rack Implementation

4

Questions


Overview

- Next Generation OCP 2S server
- DDR4
- Single/Double Sided
- Front and Rear I/O
- OpenBMC
- Open Rack V2
- Serviceability
- Sustainability(ROHS2.0/Halogen Free)
- Optimized Thermal Solution



Backward Compatibility

MB v3.0 Leopard



Add NIC Conn B to support up to 100G NIC

Add NIC Conn C to support intel KR Mezz


On-board mSATA=> M.2 conn.

PCIe x24 slot => x32 slot

BMC secure flash added

Next Gen CPU

MB v4.0 Tioga Pass




Same 6.5"x20" Form Factor

AirMax PWR + X8 => PWR*2 + X16

What's New – Double Sided SMT DIMM

- SMT DIMM socket on both PCB sides
- Maximum memory configuration
- Serviceability




What's New – X32 PCIe

- MB Connector
 - Samtec HSEC8 200-pin
 - Samtec HSEC8 60-pin
- Riser cards
 - SS 2 slot
 - SS 3 slot riser
 - DS 2 slot riser
- Features
 - JTAG support
 - Host SMBUS to slot
 - USB2.0 to slot
 - Current monitor for each slot




What's New – NIC Conn. B

- Conn. A, 120pin, PCIe Gen3 X8
- Conn. B, 80pin, PCIe Gen3 X8
- A+B to x16, support up to 100G NIC
- Conn B can not be used independently




What's New – NIC Conn. C

- 64 Pin Conn. C, up to 4x KR
- Used independently



What's New – OCP Debug Card w/ LCD


- ❑ Remapped USB3 Interface
- ❑ User Friendly
- ❑ Console(wired/wireless)
- ❑ 7 segment Post code
- ❑ LCD content
- ❑ Allow customize BMC to show debug message by needs



What's New – VGA Port


- Customized VGA port
- Pin definition

	Sig
1	RED
2	RED_RTN(GND)
3	Green
4	Green_RTN (GND)
5	Blue
6	Blue_RTN (GND)
7	V-Sync
8	GND(Vsync-DDC)
9	H-Sync
10	GND(H-sync)
11	SDA
11	SCL
13	PWR



Fast System Throttling


- CPU Dynamic Load Characteristic
- Reduce total Peak power duration
- Reduce system PSU size/cost
- Trigger sources
 - OC
 - UV
 - HSC Timer
- SEL log all sources



What's New – OpenBMC

- Facebook Open Source BMC F/W
- 1st time introduced to 2S server
- Open sourced through github

	3rd Party BMC	OpenBMC
Source Code	• Proprietary, Closed	• Designed by Facebook • Open Source @ Github
User Interface	• <code>ipmitool</code> - Raw Bytes	• SSH to BMC Linux Shell • REST API - JSON objects
Security	• RMCP+ over UDP • Known vulnerabilities	• Secure Shell (ssh) • http(s)
Authentication/Authorization	• BMC-centric username/password database	• Centralized certificate based
Multi-Node Management	• Need to virtualize BMC with multiple IP addresses	• Supported Natively with node/slot number




OpenBMC

@<https://github.com/facebook/openbmc>

What's New – BMC TPM


- I₂C Interface
- Combined with BIOS TPM

SPI_TPM_CLK	1	7	P3V3_STBY
SPI_TPM_PLTRST_N	2	8	Module_PRSNT_N/I ₂ C_T PM_RST_N
SPI_TPM_MOSI	3	9	SPI_TPM_IRQ_N
SPI_TPM_MISO	4	10	SMB CLK
SPI_TPM_CS_N	5	11	GND
I ₂ C_TPM_DAT	6		
I ₂ C_TPM for BMC			
SPI_TPM for BIOS			



What's New – Verified Boot BMC

- Build a secure BMC foundation
 - BMC chip not support CPU-based root-of-trust
 - Build with ROM based root-of-trust
- HW Support
 - Secure flash, w/Pkpub
 - BMC FW Flash, signed w/Pkpri
 - Glue Logic to reset BMC+TPM
- Verified boot Flow in U-boot





Thermal Design

- Thermal kit
 - Fan redundancy at 95°F ambient
- Thermal performance
 - Fan power is 1~2.5% of combined CPU power
 - Maximized DeltaT for thermal efficiency

Chassis Type	Single Sided			Double Sided		
Inlet Temperature (°F)	77	86	95	77	86	95
DeltaT (°F)	41.2	33.4	27.1	36.3	30.6	26.1
CFM/W	0.079	0.097	0.120	0.089	0.106	0.124

ORv2 Implementation

- 2x Zones per Rack
- 8x Chassis per Zone
- 3x Sleds per Chassis
- up to 48x sleds per rack
- 12V Medusa Cable(1->3 split)
- Pressfit Cable



Questions?



OPEN
Compute Project

