OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.

RunBMC - A Modular BMC Mezzanine Card

BUV - Bring Up Vehicle For BMC Mezzanine

Eric Shobe & Jared Mednick

Hardware Engineer - Salesforce

RunBMC

- Modular BMC Mezzanine

Eric Shobe & Jared Mednick, HW at Salesforce
eshobe@salesforce.com, jmednick@salesforce.com

Speaker Intro

Eric Shobe

Hardware Engineer
eshobe@salesforce.com

Jared Mednick

Hardware Engineer
jmednick@salesforce.com

Overview

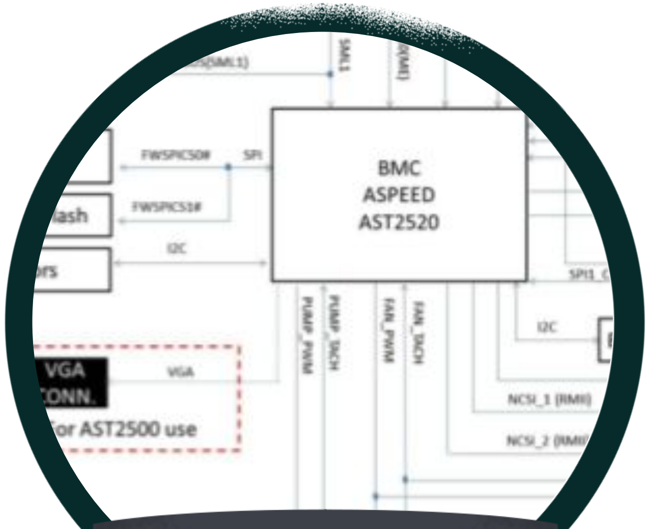

Problem Statement

RunBMC

Status

Baseboard Management Controller

Problem Statement

What is it

- BMCs (baseboard management controller) are typically specialized microcontrollers used to manage server platforms.

Problems

- OEM/ODMs treat this as a Black Box (proprietary/closed)
- Inconsistencies between each vendor! (i.e. iLo vs iDrac vs ILOM vs IMM vs RMM)
- Proprietary means harder to debug, harder to iterate, implementation of custom behaviours is challenging, etc)

Baseboard Management Controller - Software

Background

Software Stack

- Different software stacks in a BMC (Serial-Over-Lan, Sensors, Power Control, FRU’s, NC-SI, LAN Software, IPMI, USB, SNMP traps, Fan Speed/Control, Event Logs, etc)

Benefits of control and standardization, some examples

- Cost Savings to control your code
- Sensor Data
- Leverage your base code
- Security
- Increase of openBMC contributions

Baseboard Management Controller - Hardware

OCP and ODM systems (and OEMs) all share a somewhat common subsystem for the BMC. However - all are slightly different at the hardware level. This sees duplication of design effort when designing MBs for servers!

Typical Hardware Feature set (see next slide for implementations):

- AST25X0: ARM1176JZF-S ARM Processor (800Mhz)
  Co-processor (200Mhz)
- DDR4 (512MB) , PCIe Gen2
- 2x SPI NOR Flash for BMC FW (32MB)
- 1x SPI bus (Communication w/ CPLD or PCH)
- 14x I2C bus (FRU, sensors), 3x UARTs (OOB access), 2x USB 2.0
- JTAG/LPC, GPIOs/TACH/PWM, VGA
- RGMII for 1GbT Ethernet, RMII for NC-SI or 100Mb

Baseboard Management Controller

Platform Analysis

We did an in-depth analysis of currently available OCP Server platforms to compare BMC implementations

TIAGO

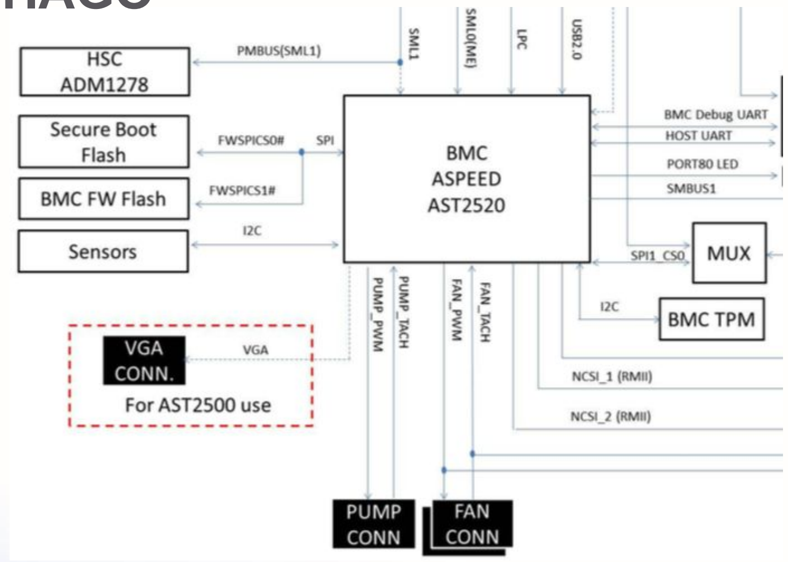

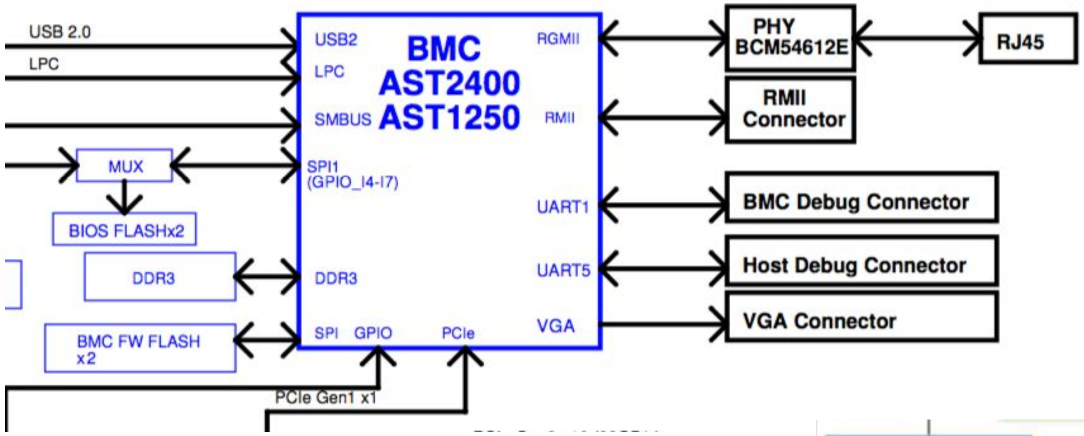

OLYMPUS

ZAIUS

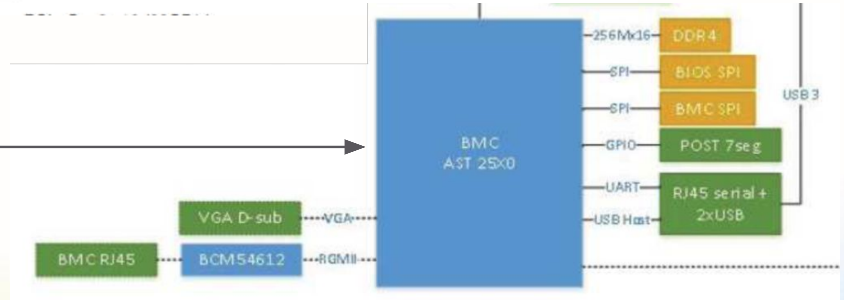

BaseBoard Management Controller - Platform Analysis

	GPIOs	IRQs	Total (GPIOs/IRQs/Others)
Tioga Pass	58	32	103
Zaius	71	20	98
Olympus	92	8	111

Much of this signal functionality is typically shared with the PCH (for example, control of Host voltage supply). Not on Zaius!

Wedge BMC Analysis

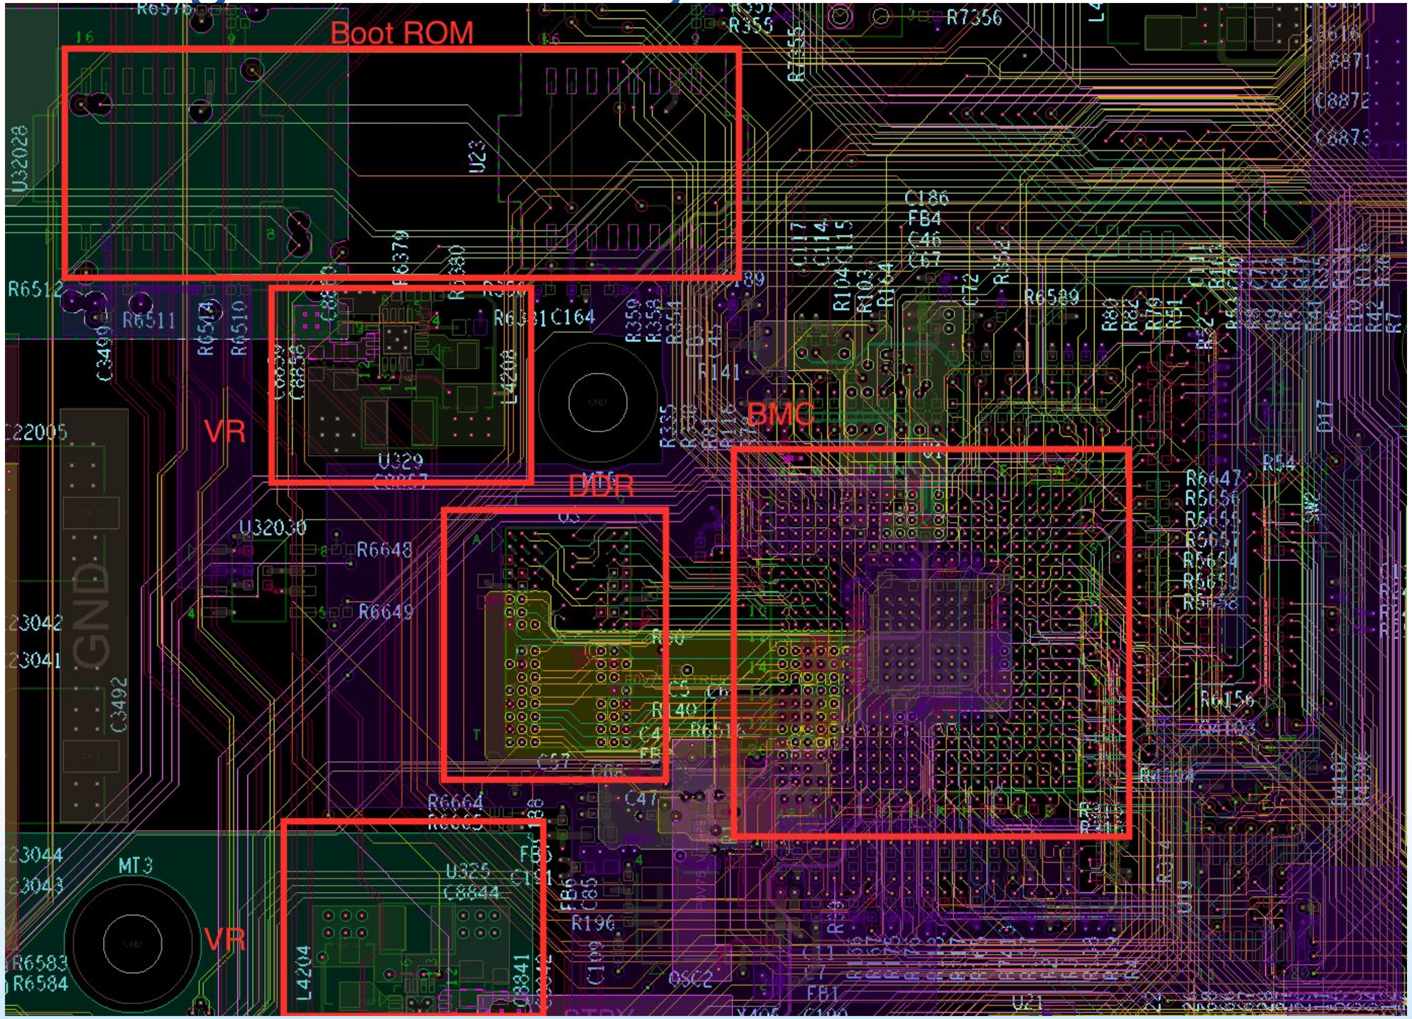

Tioga Pass BMC Analysis

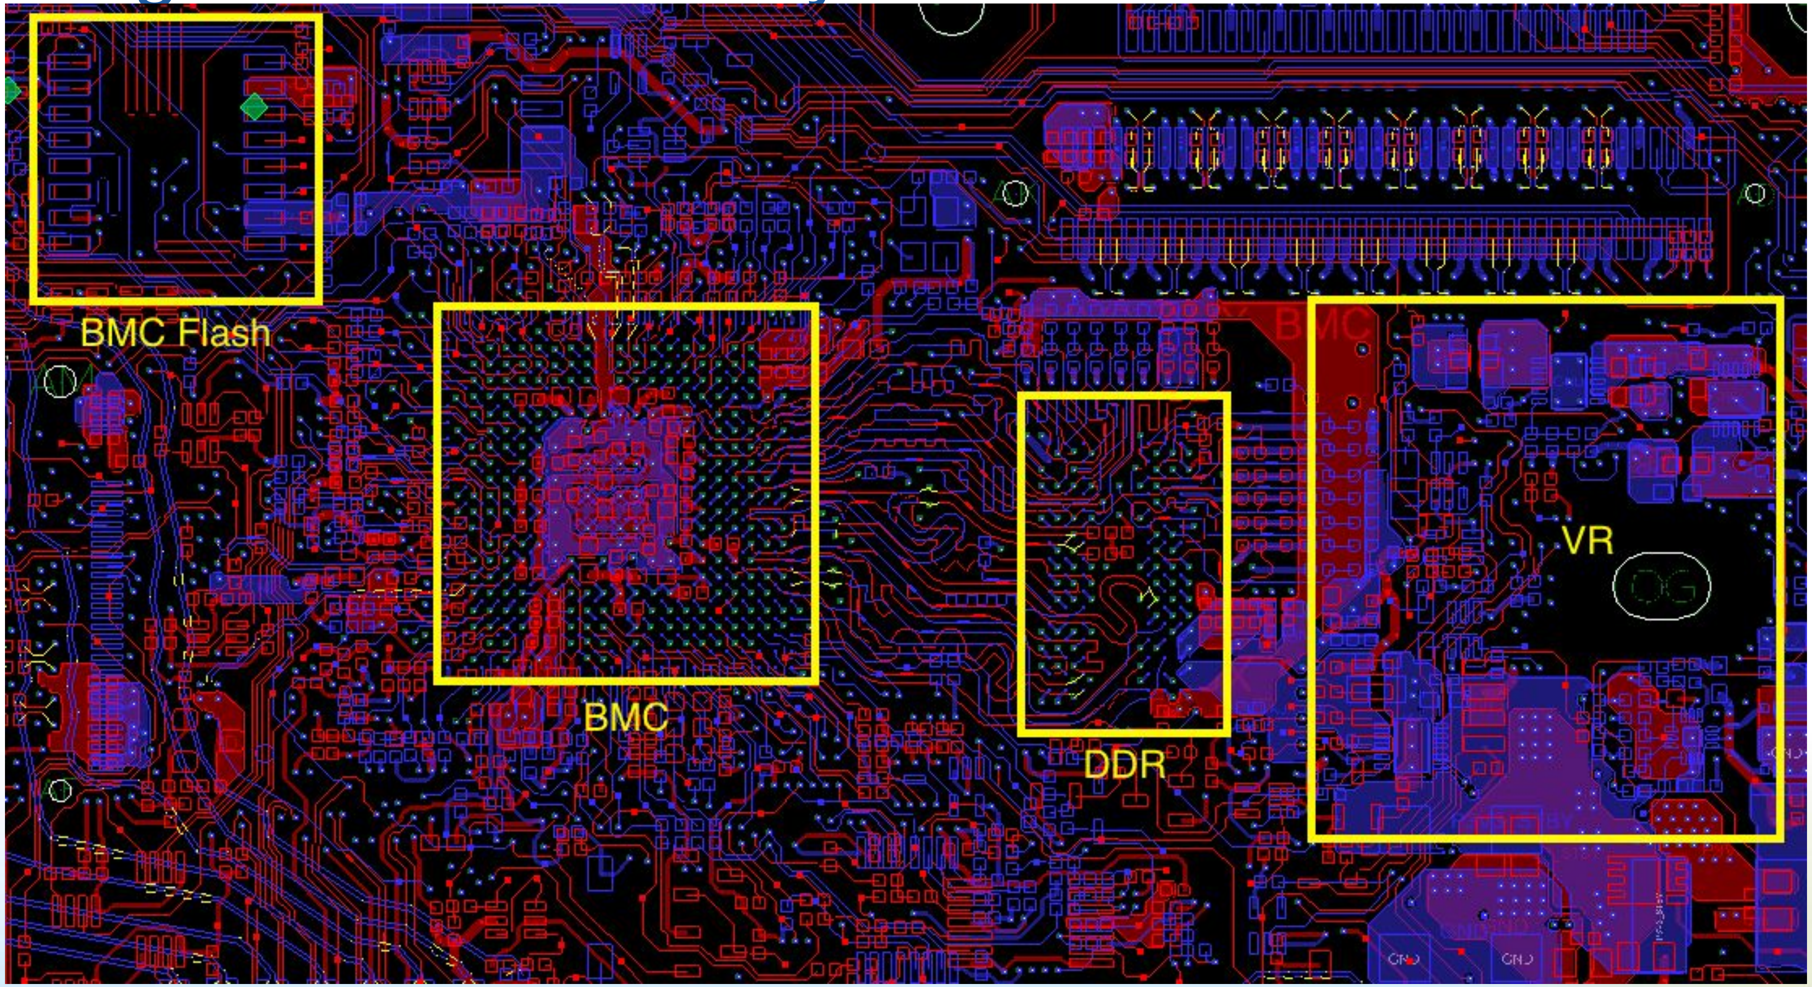

Zaius BMC Analysis

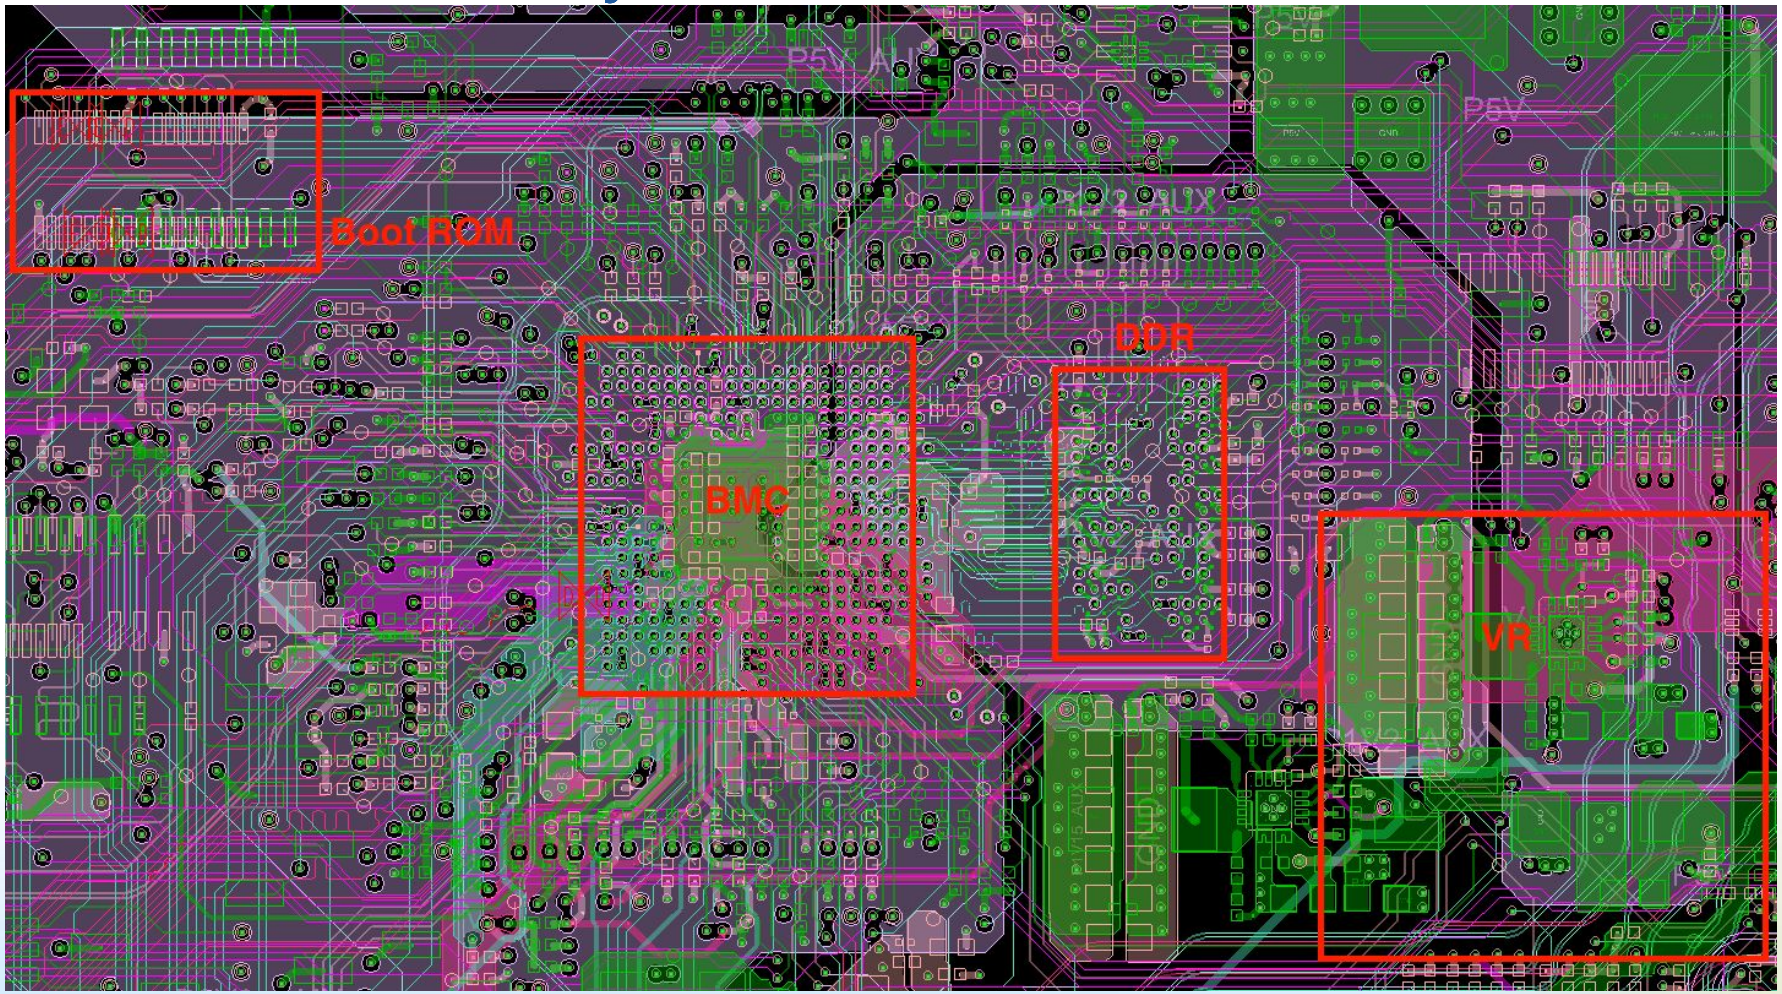

Olympus BMC Analysis

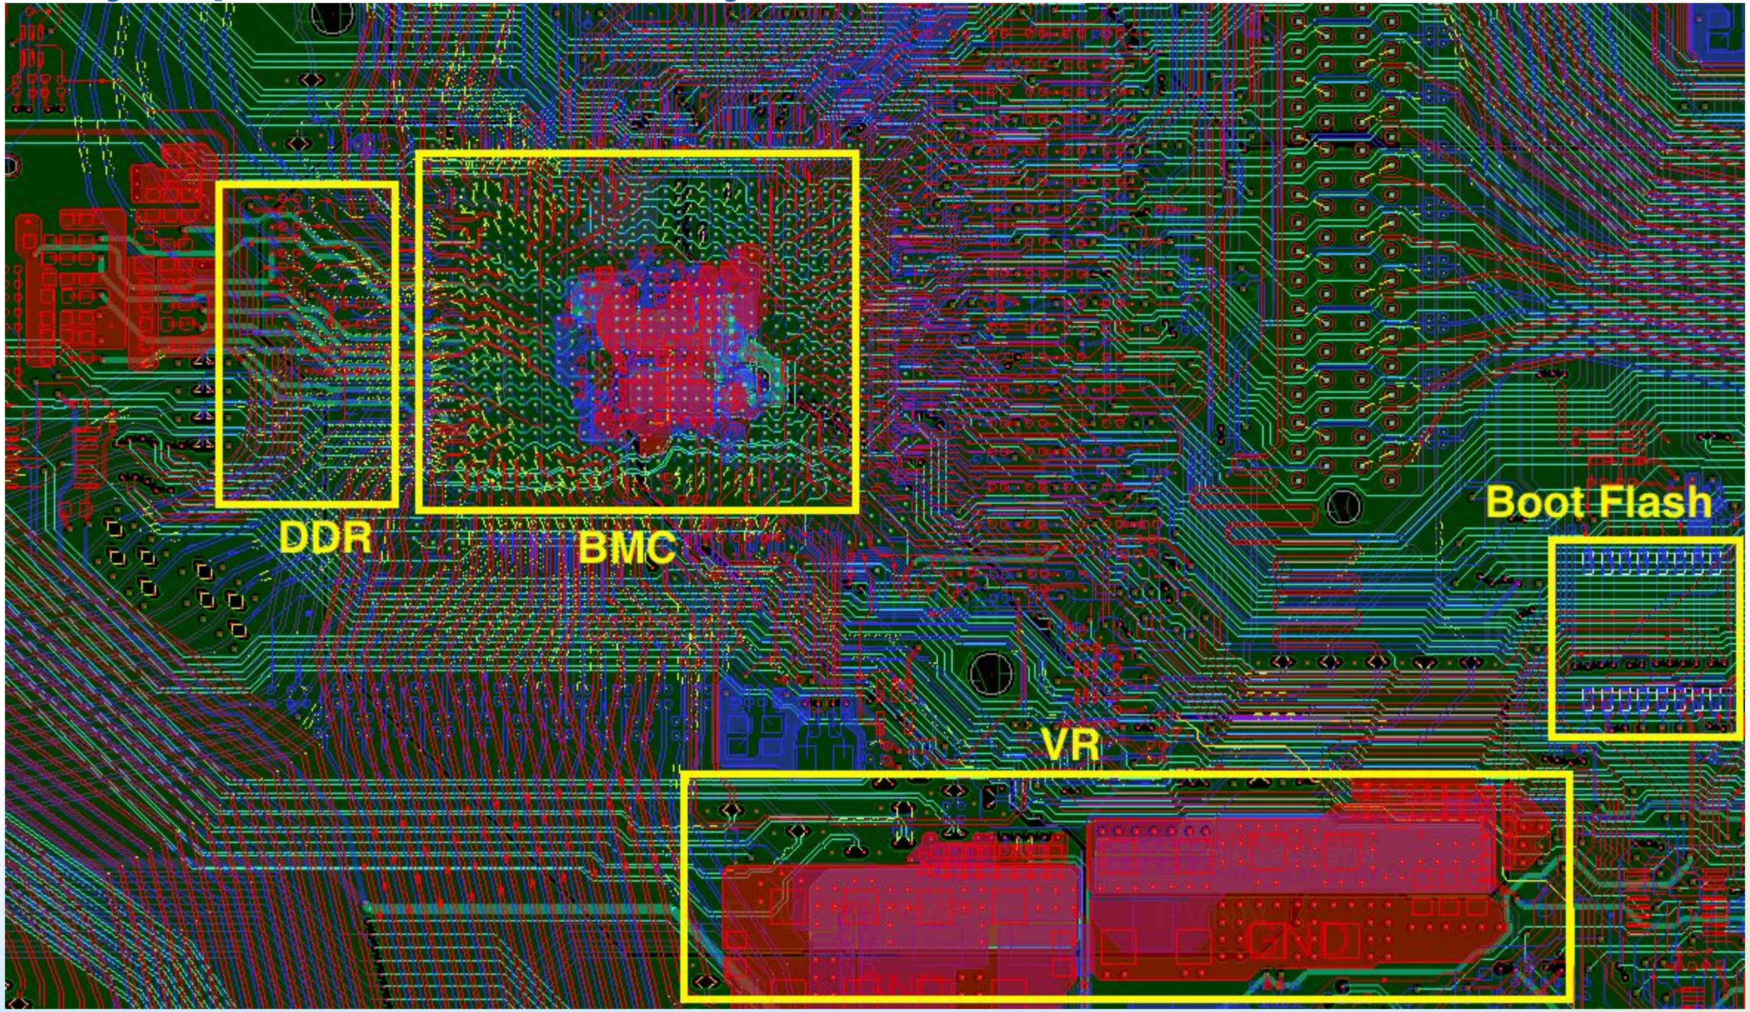

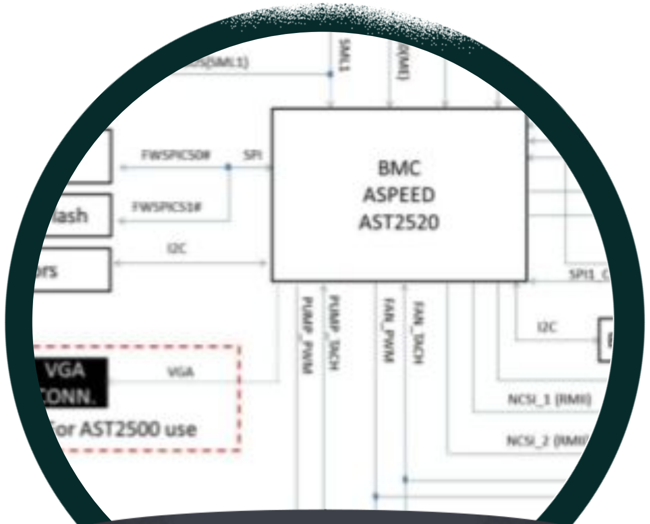

Problem Statement

RunBMC

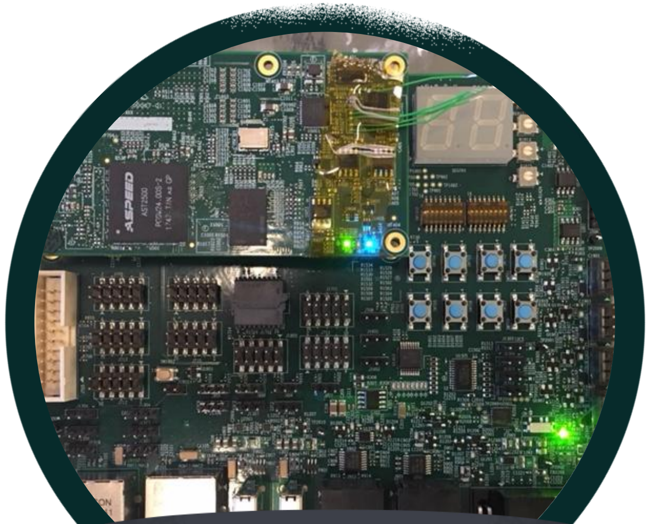

Status

RunBMC

Our Solution

- Modular Design
- Standardize Pinout / Connector
- Standardize the HW footprint

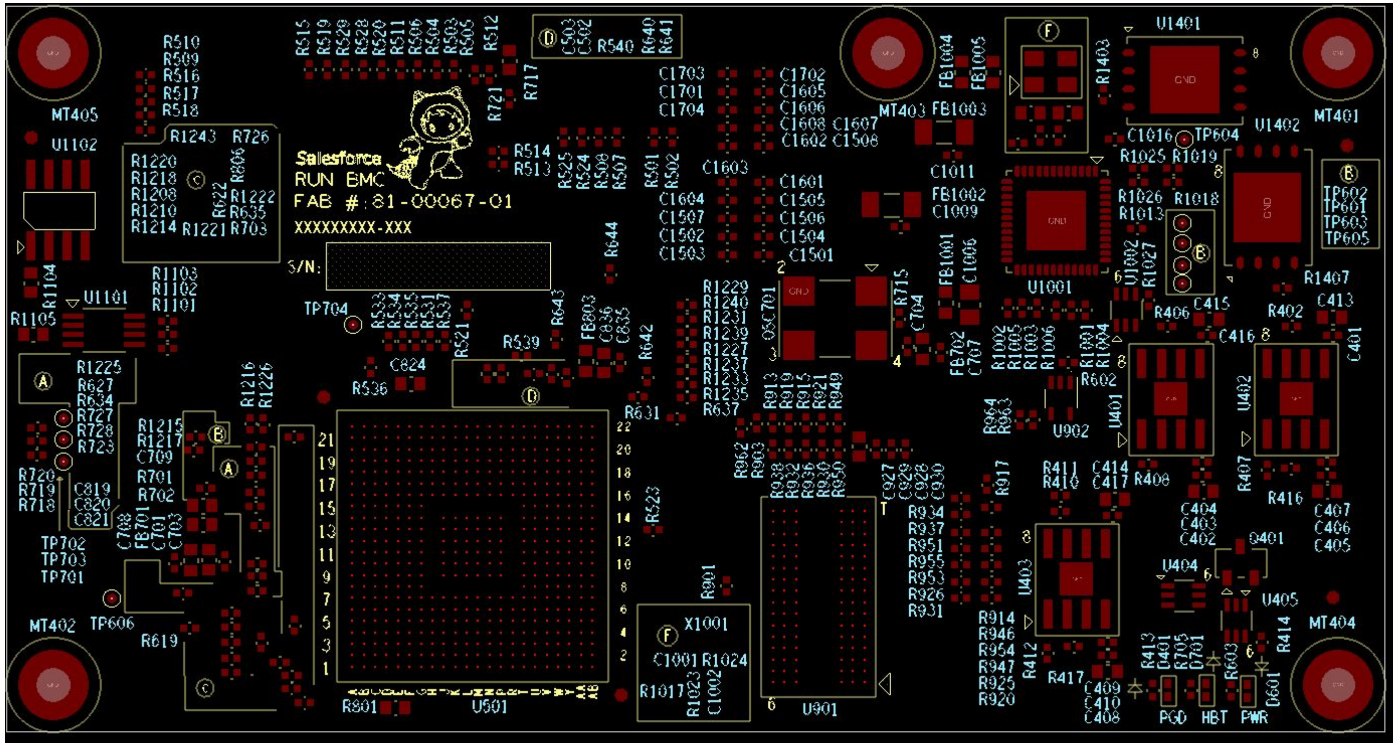

RunBMC - Connector Pinout

	A	B	C	D	E	F	G	H	I	J
1	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V
2	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V	3.3V
3	GND	GND	GND	GND	GND	GND	GND	GND	GND	GND
4	ADC11	PHY_LED1	GND	GND	GND	GND	GND	GND	PEREFCLKP	GND
5	ADC12	GND	LPC_CLK	GPIOM7	PWM0	GPIOH5	I2C1_SCL	GND	PEREFCLKN	PHY_MDI_P0
6	ADC13	PHY_LED2	LPC_FRAME_N	GPIOF0	PWM1	GPIOH6	I2C1_SDA	UART_TXD4	GND	PHY_MDI_N0
7	GND	GPIOL2	LPC_SERIRQ_N	GPIOF1	PWM2	GPIOH7	GND	UART_RXD4	PERXP	GND
8	GPIOQ4	GPIOL3	LPC_AD0	GND	GND	GPIOJ0	I2C5_SCL	GND	PERXN	PHY_MDI_P1
9	PECIVDD	GPIOL4	LPC_AD1	GPIOF2	PWM3	GND	I2C5_SDA	UART_TXD1	GND	PHY_MDI_N1
10	GPIOQ5	GPIOE0	LPC_AD2	GPIOF3	GND	SYSCK	GND	UART_RXD1	PERST_N	GND
11	GPIOQ6	GPIOE1	LPC_AD3	GPIOF4	PWM4	SYSMOSI	I2C11_SCL	GND	PECI	PHY_MDI_P2
12	GPIOQ7	GND	GND	GPIOH0	PWM5	SYSMISO	I2C11_SDA	UART_RTS4	GPIOD6	PHY_MDI_N2
13	GND	GPIOE2	SPI2_CS1_N	GPIOH1	PWM6	SYSCS_N	GND	UART_RTS1	GND	GND
14	ADC14	TACH14	JTAG_TDI	GPIOH2	PWM7	GND	I2C2_SCL	GND	JTAG_TRST	PHY_MDI_P3
15	ADC15	TACH15	JTAG_TMS	GPIOH3	SPI2_CS0_N	GPIOJ1	I2C2_SDA	UART_TXD3	MDIO1	PHY_MDI_N3
16	GND	GPIOE3	GND	GPIOH4	GND	GPIOJ2	GND	UART_RXD3	GND	GND
17	SPI2_SCK	LPC_RST_N	JTAG_TCK	GND	RMII1TXEN	GPIOJ3	I2C3_SCL	GND	GPIOD7	MDC1
18	GPIOL0	GPIOE4	JTAG_TDO	DACB	RMII1TXD0	GPIOD0	I2C3_SDA	UART_TXD5	GPIOA0	ADC4
19	GND	GPIOE5	GND	DACG	RMII1TXD1	GPIOD1	GND	UART_RXD5	GND	GND
20	RMII1RCLKI	GND	GND	DACR	GPIOT4	GPIOD2	I2C4_SCL	GND	PETXP	ADC0
21	GPIOU5	PHY_LED3	TACH0	VGAHS	GPIOT5	WDTRST1	I2C4_SDA	I2C12_SDA	PETXN	ADC2
22	RMII1RXD0	GPIOM0	TACH1	VGAVS	GND	WDTRST2	GND	I2C12_SCL	GND	ADC1
23	RMII1RXD1	GPIOM1	TACH2	GND	JTAG_RTCK	GND	I2C7_SCL	GND	RST_N_CONN	GND
24	RMII1CRSDV	GPIOM2	TACH3	DDCCLK	GPIOD4	SPI1CK	I2C7_SDA	I2C10_SDA	ADC5	ADC8
25	RMII1RXER	GPIOM3	TACH4	DDCDAT	SPI2_MOSI	SPI1CS0_N	GND	I2C10_SCL	ADC6	GND
26	GPIOL1	GPIOM4	TACH5	GND	SPI2_MISO	SPI1MISO	I2C8_SCL	GND	ADC7	ADC9
27	GND	GPIOM5	TACH6	TACH10	GND	SPI1MOSI	I2C8_SDA	I2C6_SDA	ADC3	GND
28	USB2B_DP	GND	TACH7	TACH11	USB2A_DP	GND	GND	I2C6_SCL	GND	ADC10
29	USB2B_DN	GND	TACH8	TACH12	USB2A_DN	GND	I2C9_SCL	GND	GPIOA1	GND
30	GND	GPIOM6	TACH9	TACH13	GND	GPIOD3	I2C9_SDA	GPIOD5	GPIOA2	GPIOA3

RunBMC - Connector Pinout

Summary

- Example Configuration for our reference board:

Power	20
GND	80
ADC	16
PCIE	7
1 GbT	13
VGA	7
RMII	8
JTAG	6
USB	4
SPI	13
LPC	8
I2C	24
UART	10
PWM	8
TACH	16
PECI	2
GPIO	55
RESET	3
TOTAL	300

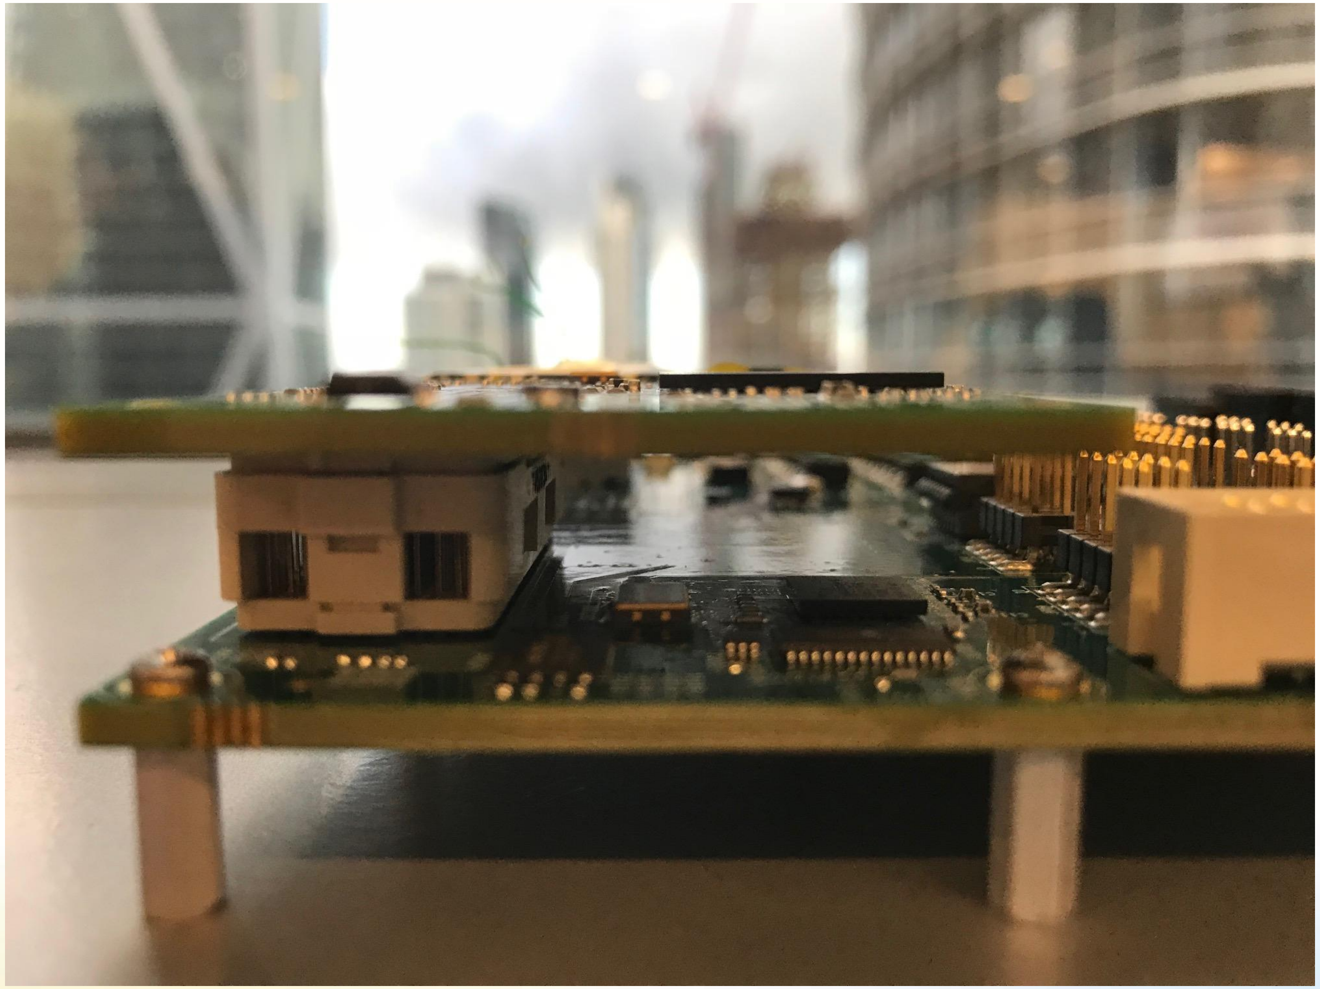

RunBMC

Challenges

- Feature Limited
- Form factor
- Adoption
- Price!

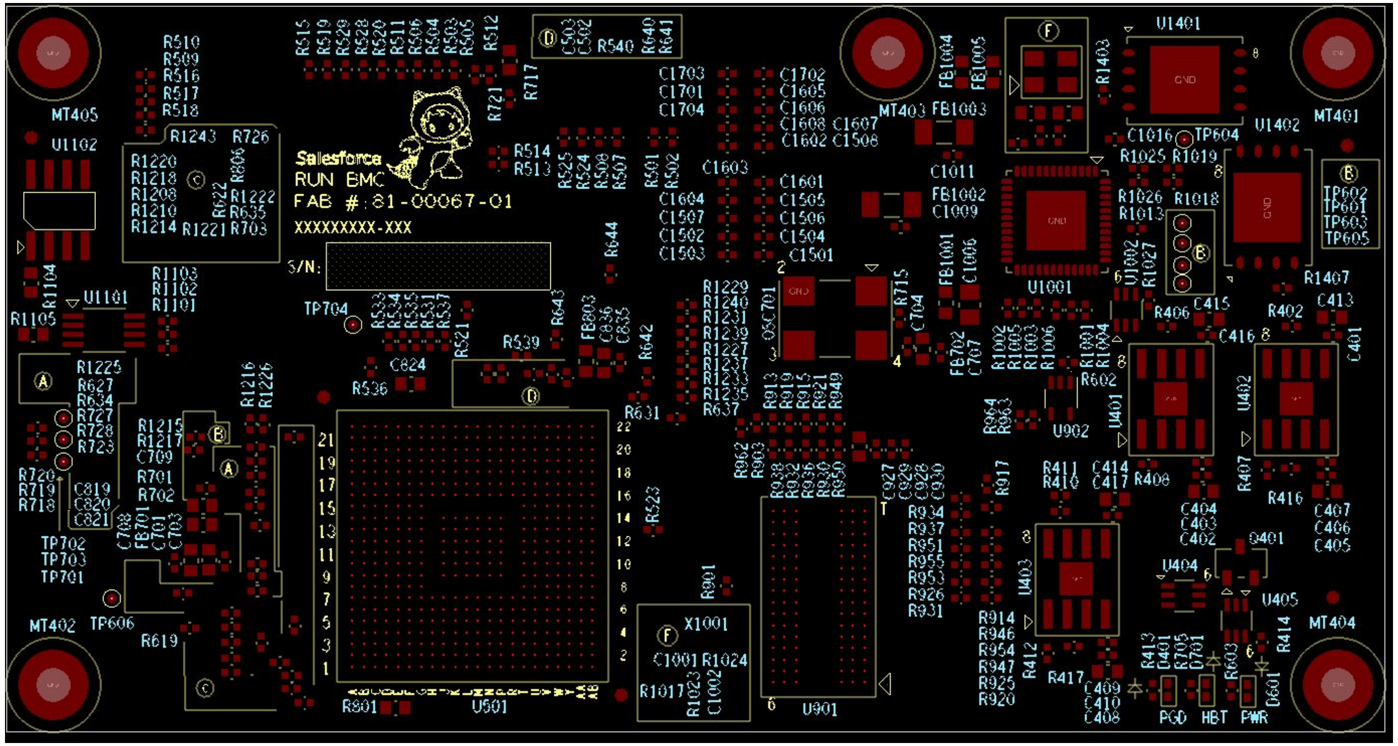

RunBMC - Reference Board

Features

- ASPEED AST2500 SOC
- DDR4 (512MB)
- 300pin .7mm pitch connector
- 1Gb Phy for RGMII (BCM54612e)
- Power Delivery standardized
- 2x SPI NOR for BMC FW
- Local I2C
- 3x UARTs, 2x USB, JTAG, LPC
- 2x SPI
- Tach/PWM
- 12x I2C
- VGA
- 1GbT and NC-SI

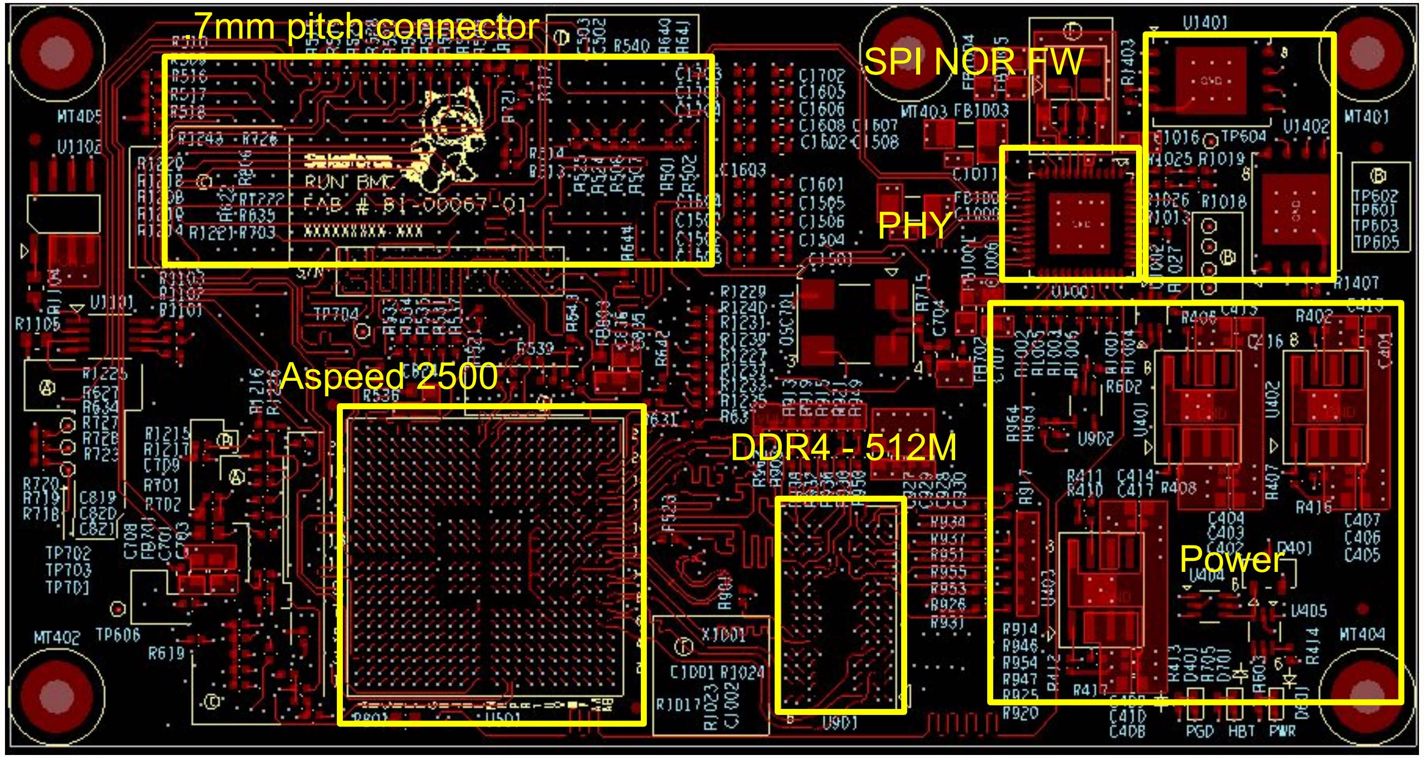

BUV - Bring Up Vehicle

Our Solution

- Bring-up board allows easy development cycles w/out hardware overhead. Access to network interfaces, UART, some basic parts for testing. Can be used at your desk, universities, etc

Features

- RMII to 10/100 PHY
- 2x RJ45 for 1Gb and 100Mb
- Micro-USB access for console and power
- LPC, JTAG headers
- 12x I2C headers
- GPIO/ADC headers
- 7 seg display
- 3x console
- JTAG/LPC
- TPM
- 2x SPI bus
- TACH/PWM
- I2C devices (TPM, Temp, Current, expander)
- VGA
- Clocks

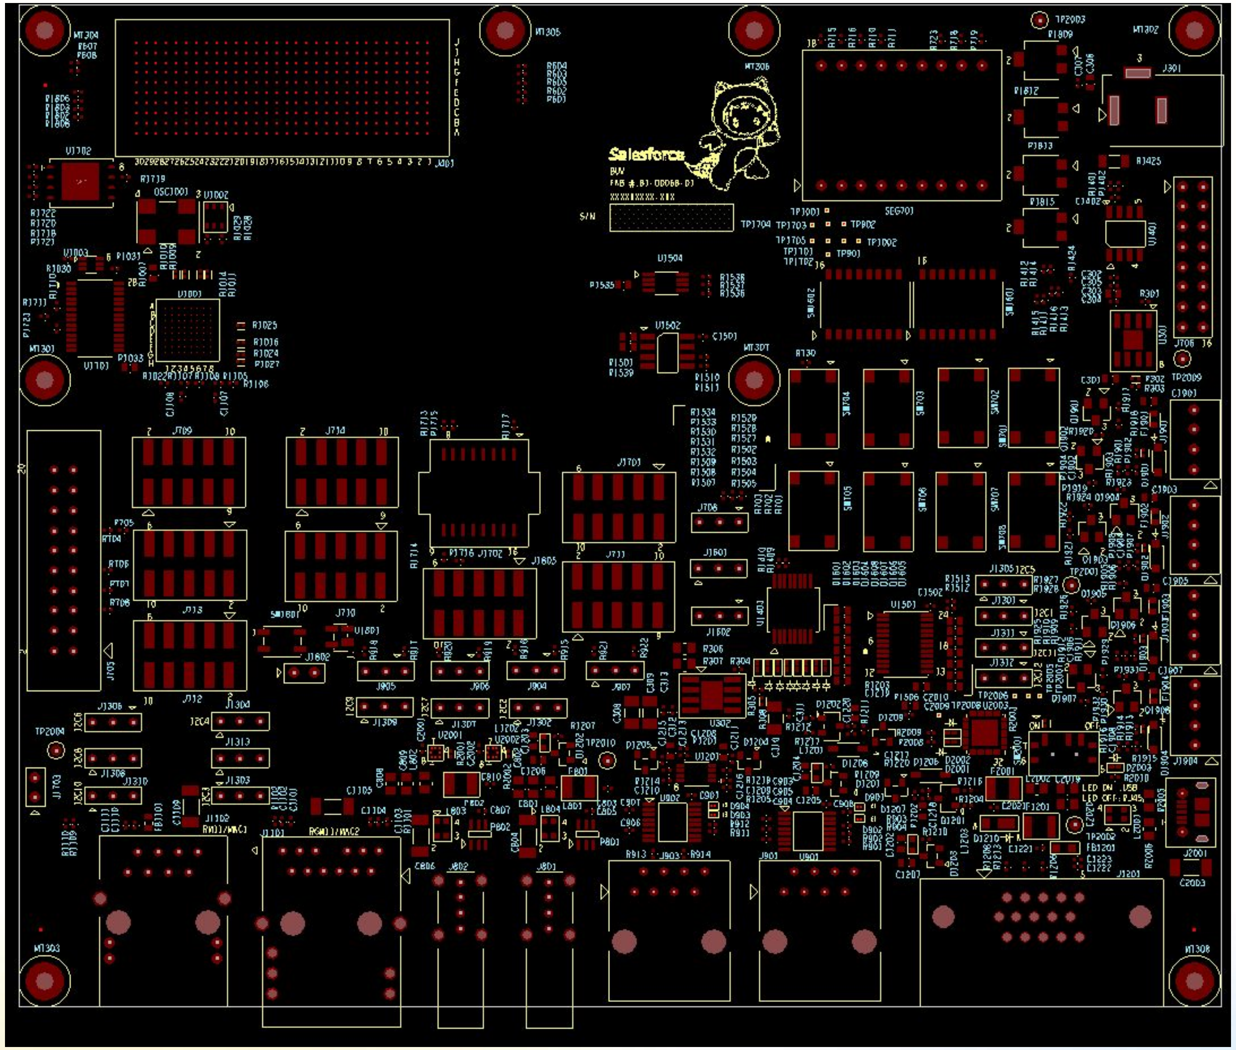

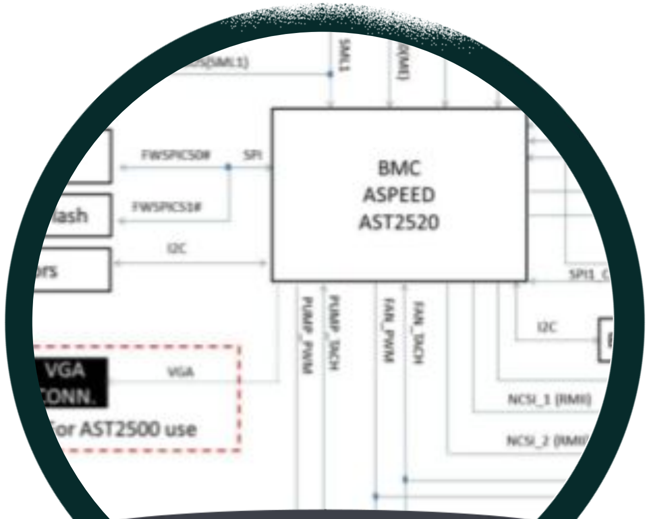

Problem Statement

RunBMC

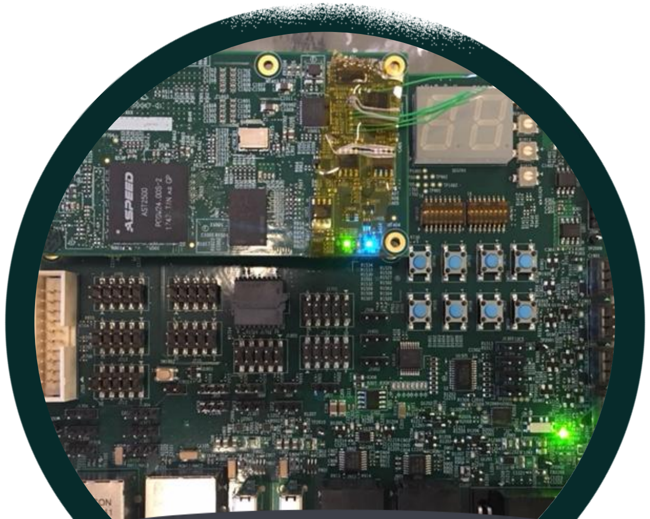

Status

RunBMC Mezzanine - Stage I | RunBMC + Bring Up Vehicle

salesforce

Goal is to create a stand-alone BMC mezzanine (RunBMC) + bring-up board. Low risk project compared to full platform port (stage II). Success is defined by complete bring-up of board, booting linux, and ssh'ing into BMC.

RunBMC mezzanine board	A re-usable daughter card containing the BMC chip (ASPEED AST2520), common subsystems (i2c, power, ddr, clocks, eeproms, flash), and a small pitch connector for any hardware platform to use (switches, servers, etc). Only the platform SW changes.
Bring-up vehicle board	RunBMC mezzanine and Bring-up board allows easy development cycles w/out hardware overhead. Access to network interfaces, UART, some basic parts for testing. Can be used at your desk, universities, etc
Re-use	RunBMC mezzanine would allow faster turnaround from a hardware development life-cycle, earlier platform software development, and higher software adoption from developers (using the development bring-up board)

RunBMC - Bring-Up Success!

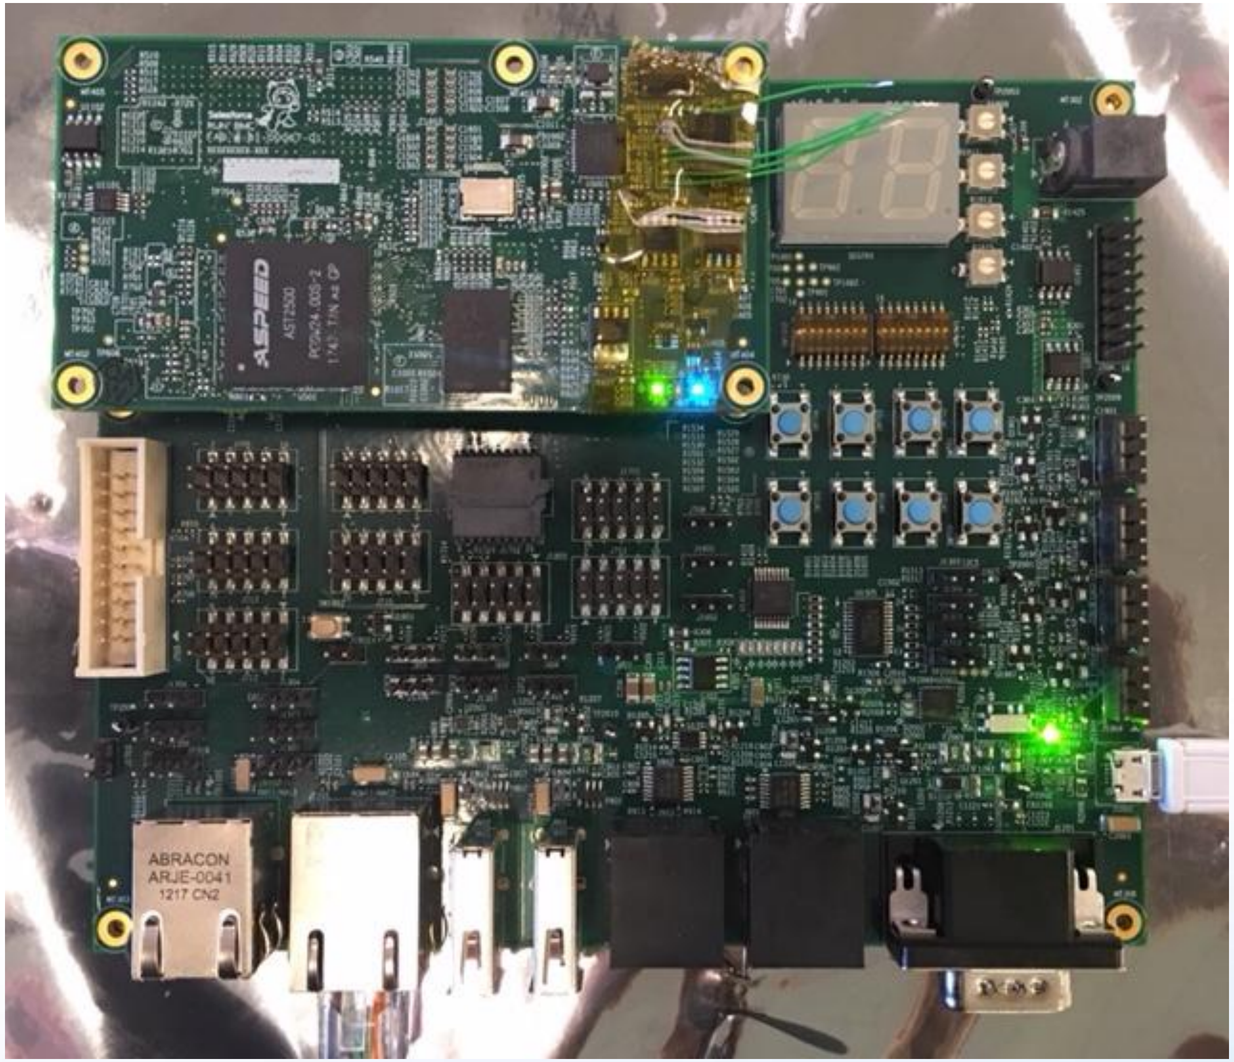


```
runBMC#
runBMC#
runBMC# bdinfo
arch_number = 0x000022B8
boot_params = 0x80000100
DRAM bank   = 0x00000000
-> start    = 0x80000000
-> size     = 0x1F000000
eth0name    = aspeednic#1
ethaddr     = 00:11:22:33:44:55
current eth = aspeednic#1
ip_addr     = 192.168.1.2
baudrate    = 115200 bps
TLB addr    = 0x9EFF0000
relocaddr   = 0x9EFAA000
reloc off   = 0x9EFAA000
irq_sp      = 0x9EB88EF0
sp start    = 0x9EB88EE0
runBMC#
Meta-Z for help | 115200 8N1 | NOR | Minicom 2.7.1 |
```

RunBMC Mezzanine - Stage II | RunBMC + Platform Port

salesforce

Integrate RunBMC mezzanine with ODM server platform of choice. Riskier port due to interactions with CPU, PCH, CPLD/FGPA subsystems, power-on, and deeper software integration. However - lots of code will be re-used.

RunBMC mezzanine board	The re-usable daughter card defined in Stage I.
ODM server platform	Our ODM platform modified (schematic/layout) to use the BMC mezzanine board.
Software Integration	Deeper integration to SFDC tools (coolan agents on the BMC), KISS (keep it simple stupid - less software bloat), etc.

Questions?


THANK YOU

OCP
SUMMIT