



OCP SUMMIT

March 20-21
2018
San Jose, CA



OPEN
Compute Project



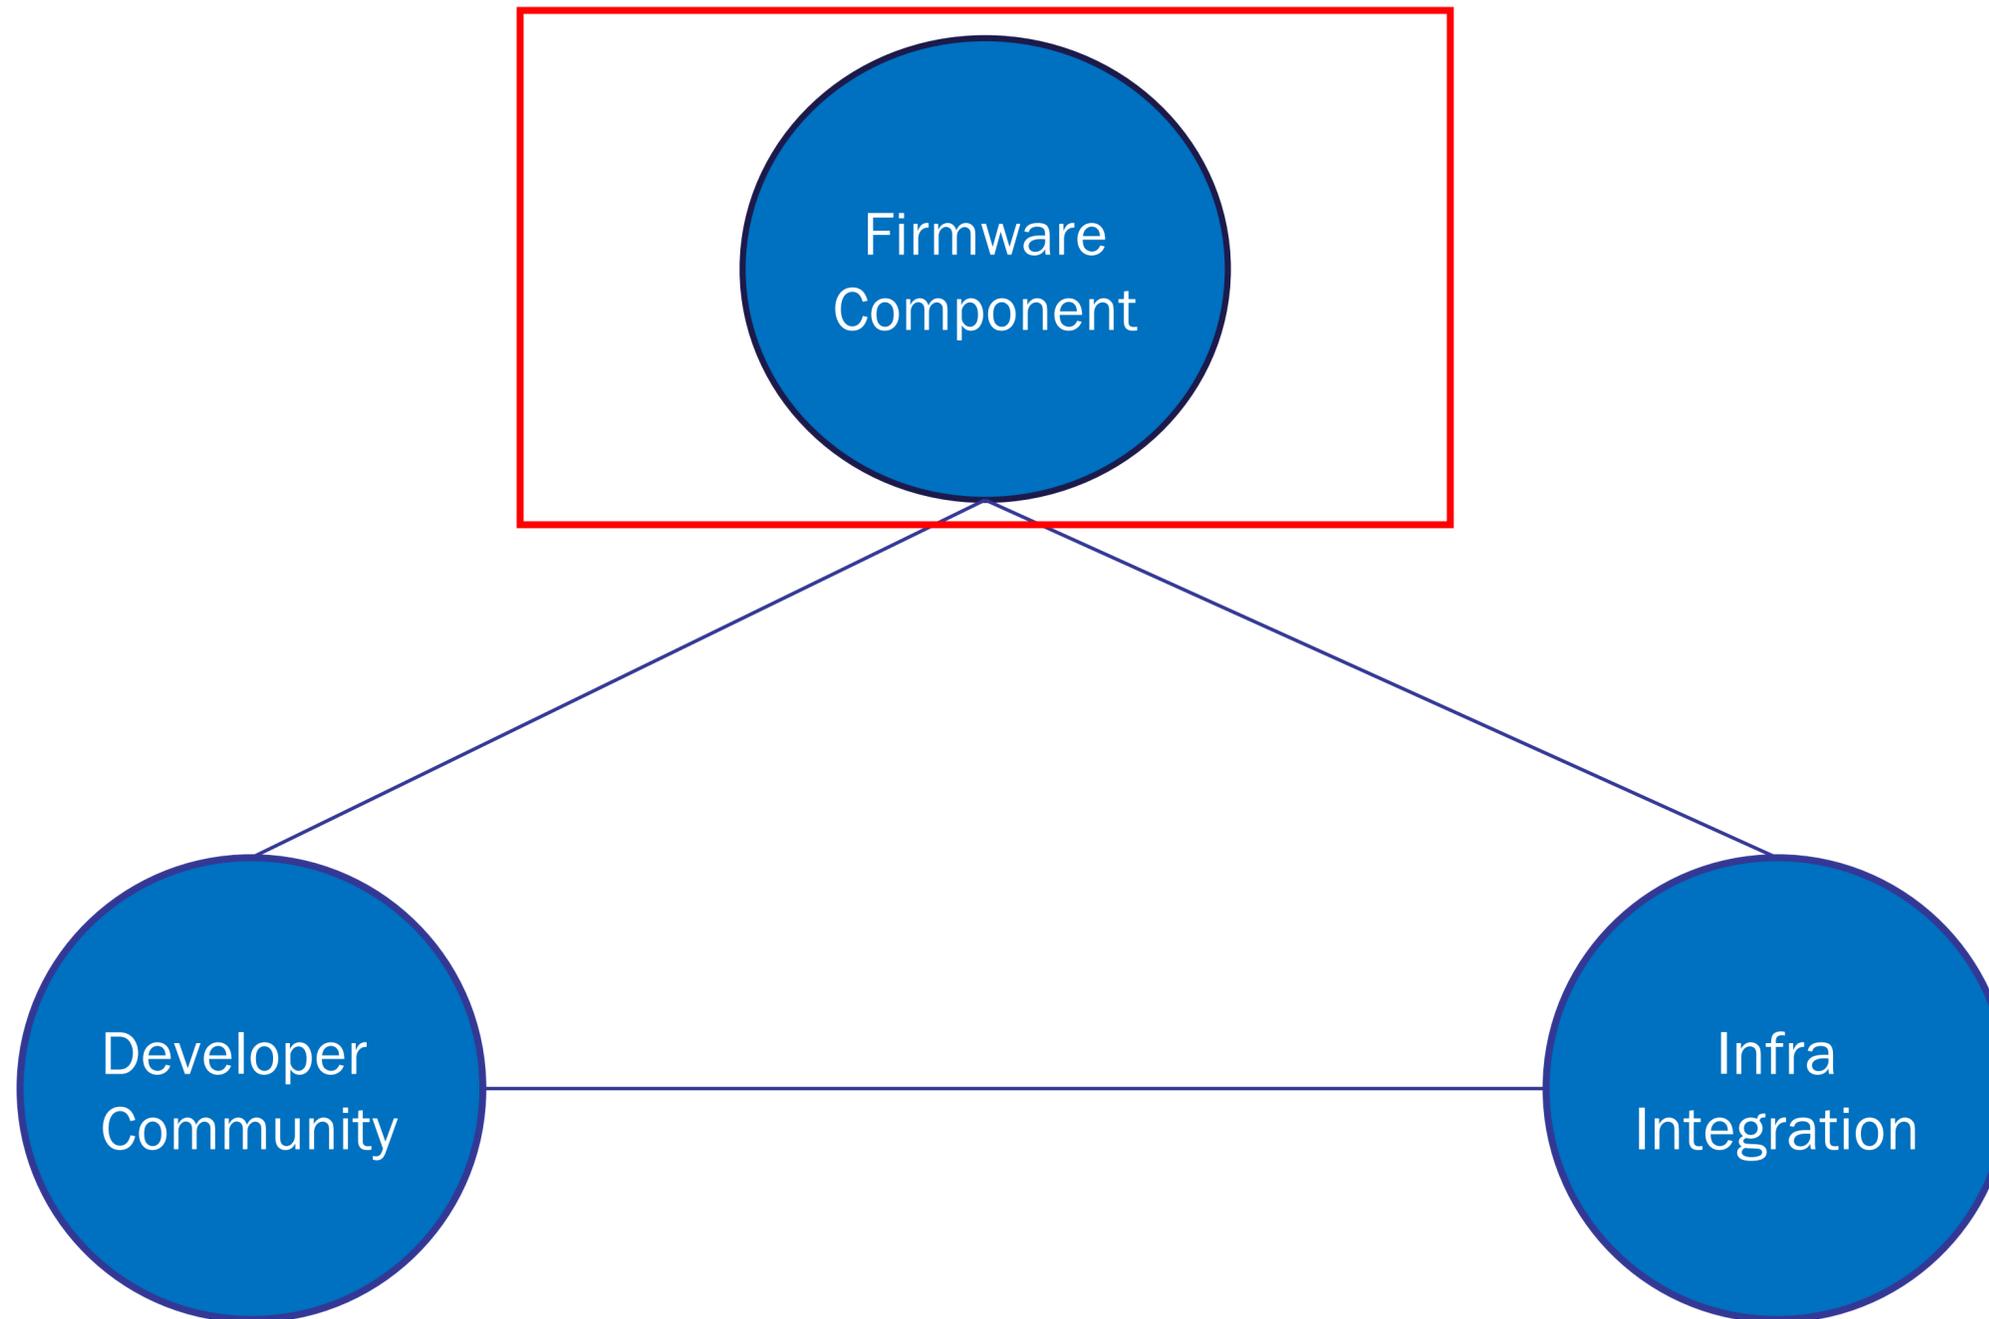
Facebook OpenBMC Updates

Sai Dasari, Christopher Covington
Facebook

OCP Summit Recap

- **OCP Summit'15**
 - Proof-Of-Concept - Wedge
- **OCP Summit'16**
 - Server Management - Yosemite
 - Storage Management - Lightning
- **OCP Summit'17**
 - New Platforms - Wedge100, Backpack, YosemiteV2, TiogaPass, Bryce Canyon
 - Updates - AST25xx, U-Boot/2016.07, Linux/4.1.15
 - Security Feature - Verified Boot POC

OCP '18 Agenda

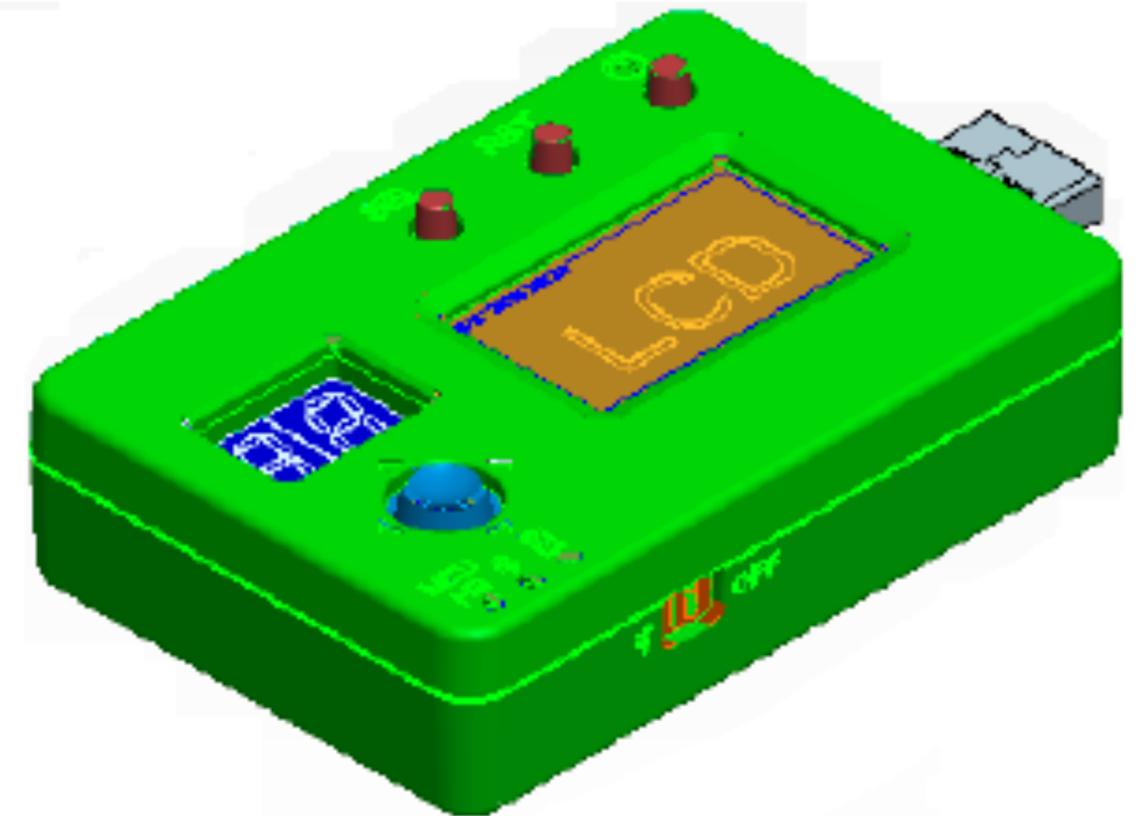
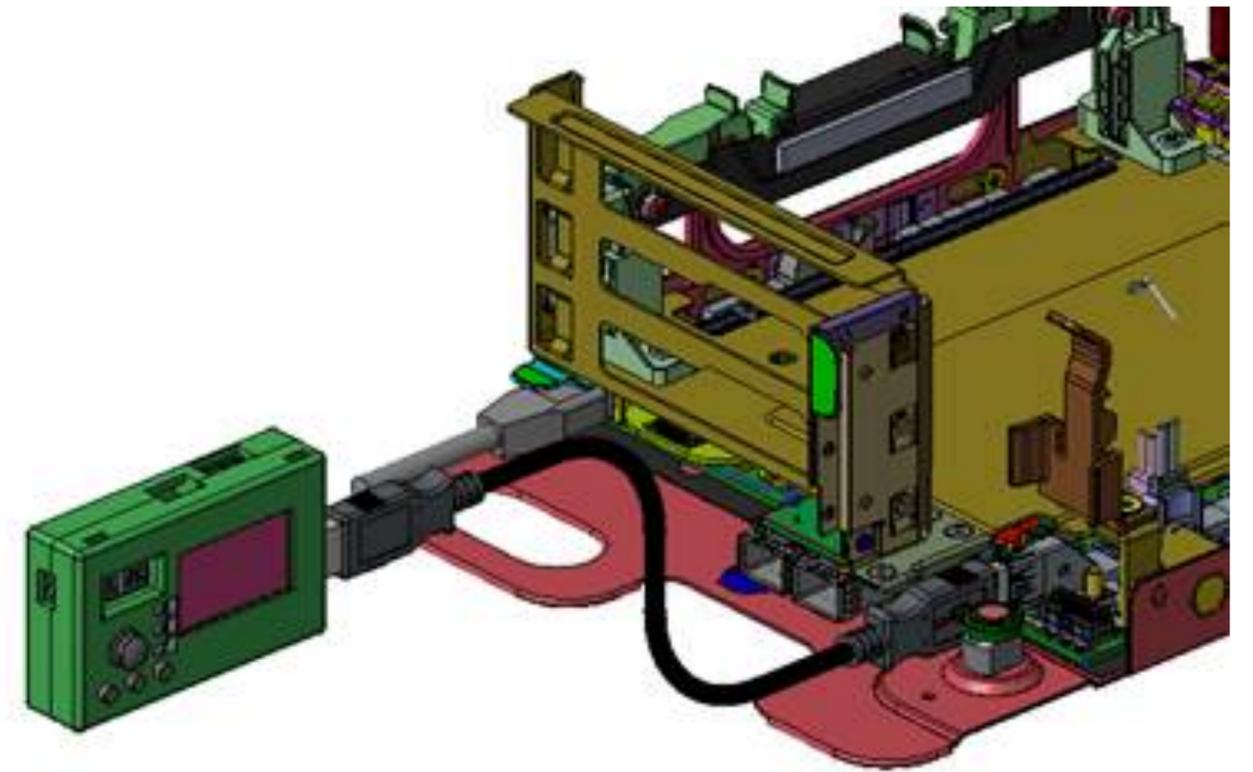


Since Last Summit

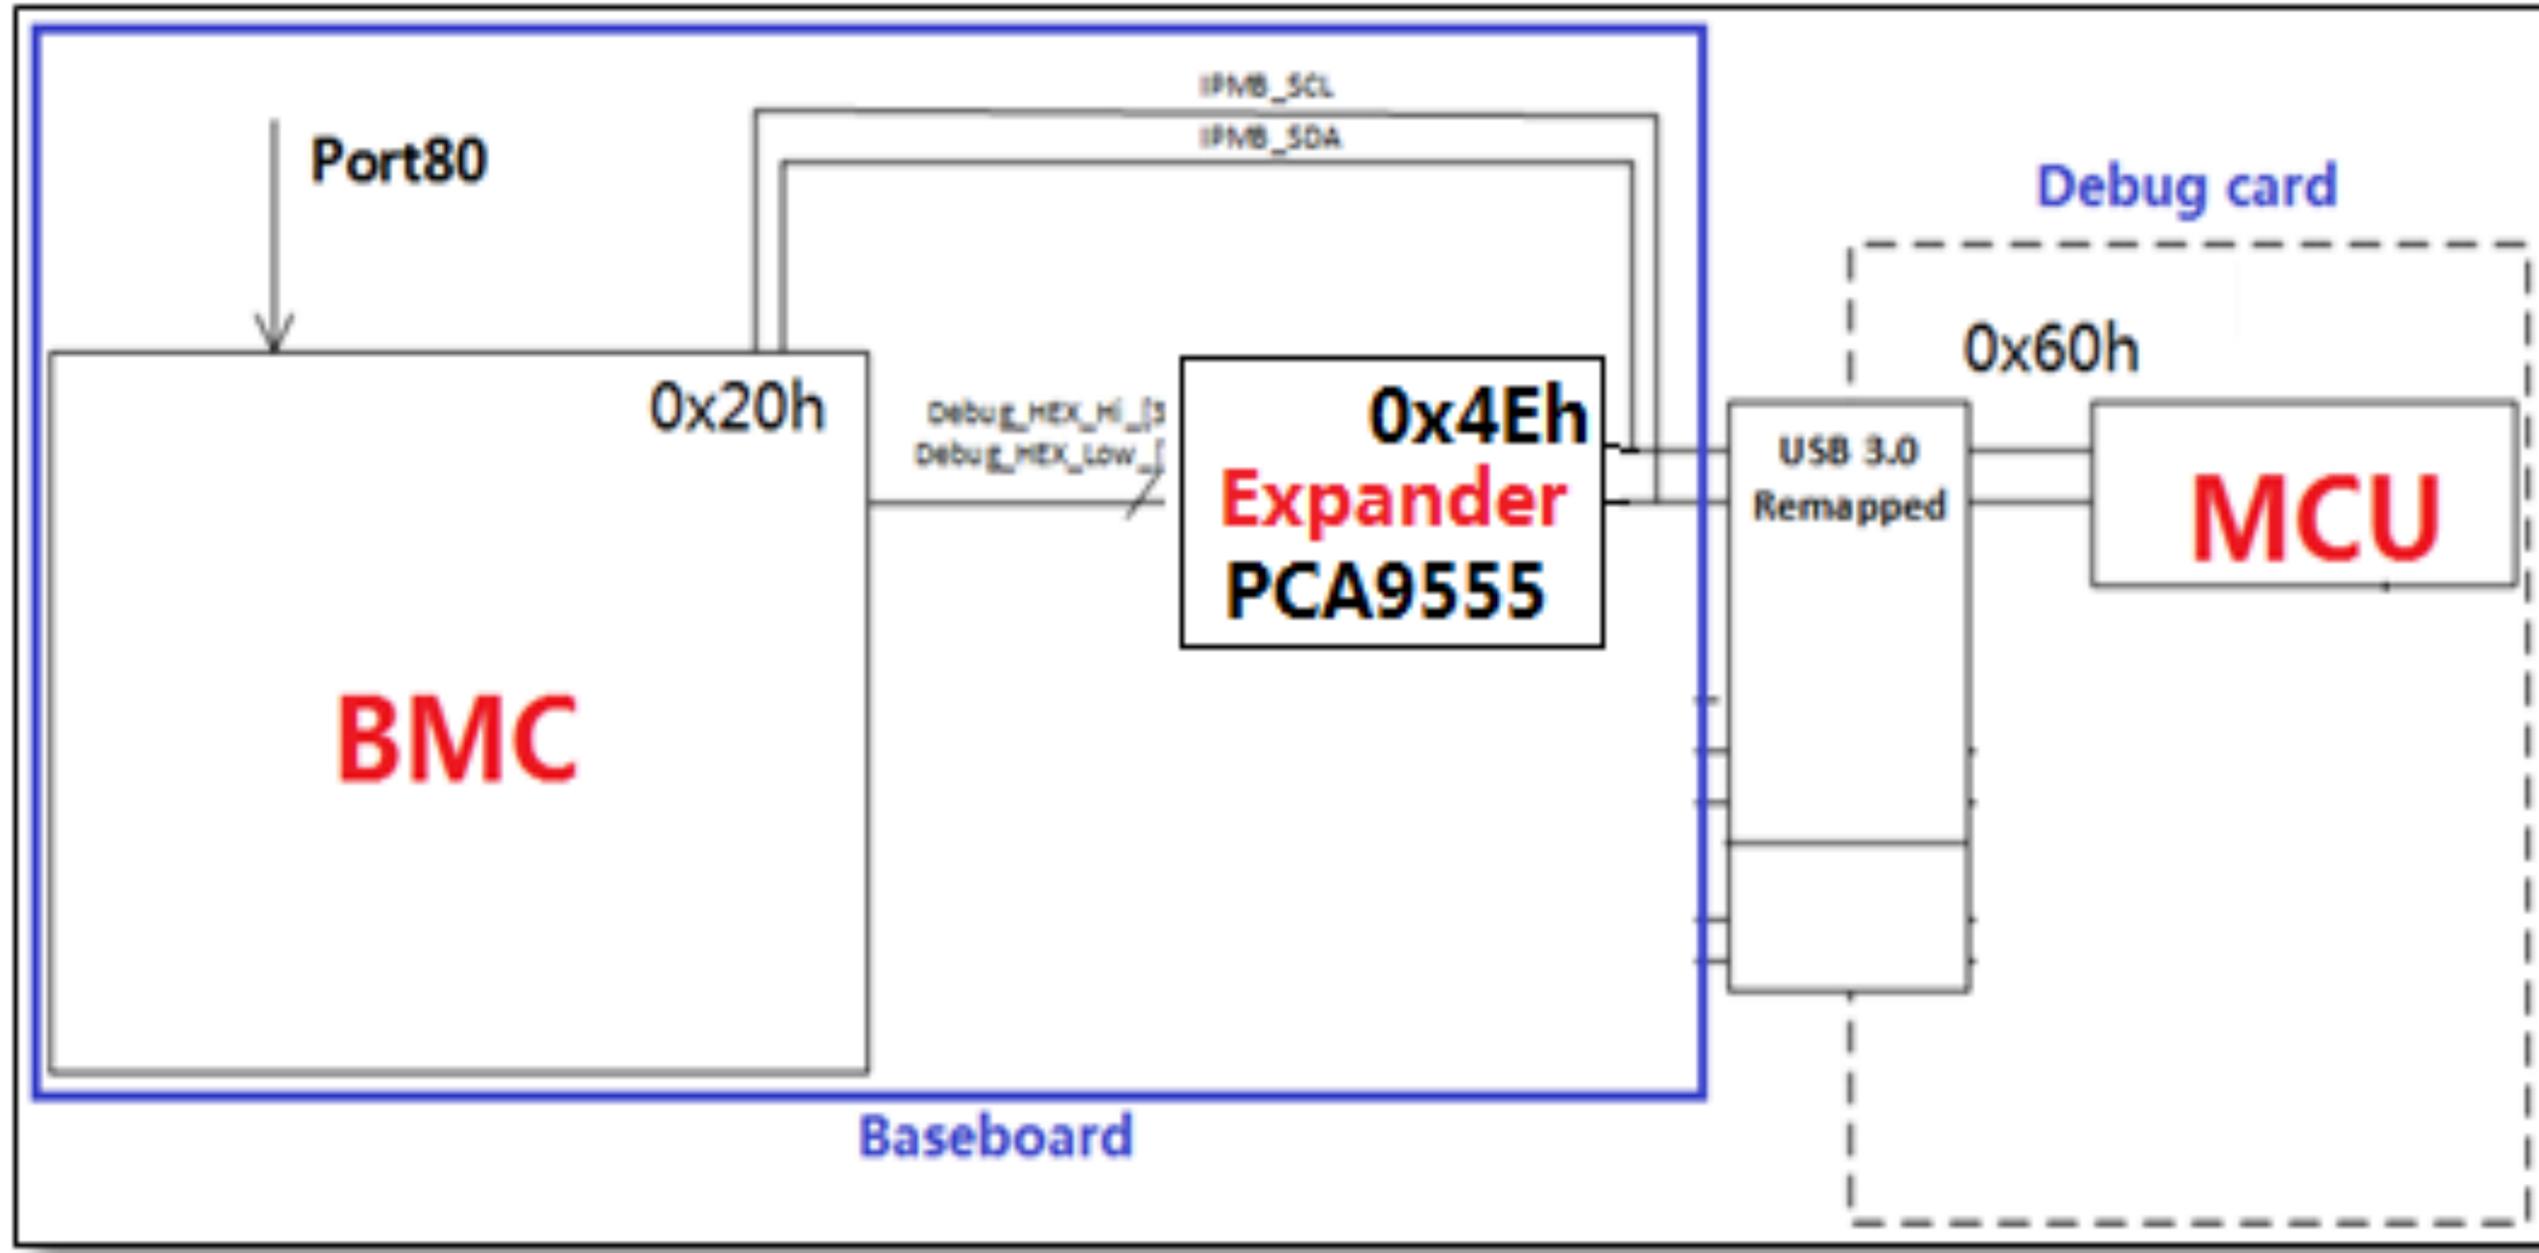
- Updates
 - Python2->Python3 Migration
 - Yocto: Krogoth->Rocko
 - Verified Boot
- Features:
 - YosemiteV2 Hot Service
 - OCP LCD Debug Card
 - At-Scale Debug

OCP LCD Debug Card

- LCD
 - System Information
 - POST Code
 - GPIO Status
 - SEL
 - Critical Sensors
- Control
 - Power
 - Reset
 - Console Selection
 - Hot Service



Host System->Debug Card Interface



Example Debug Pictures#2

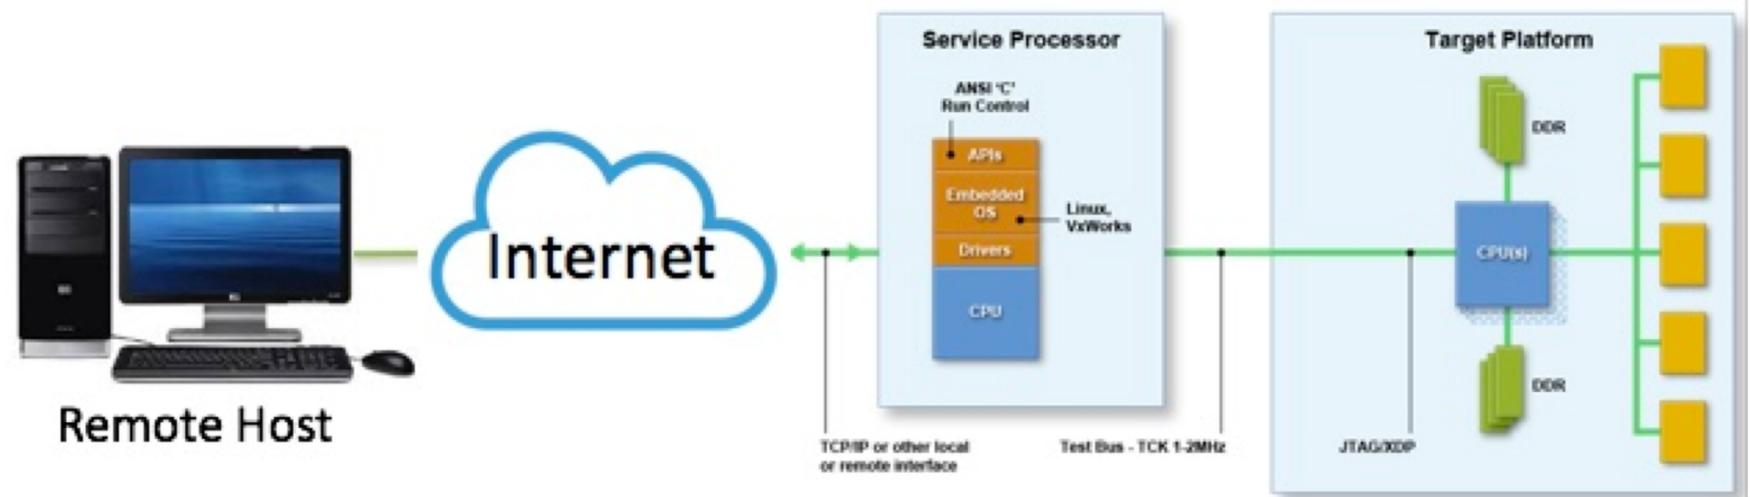
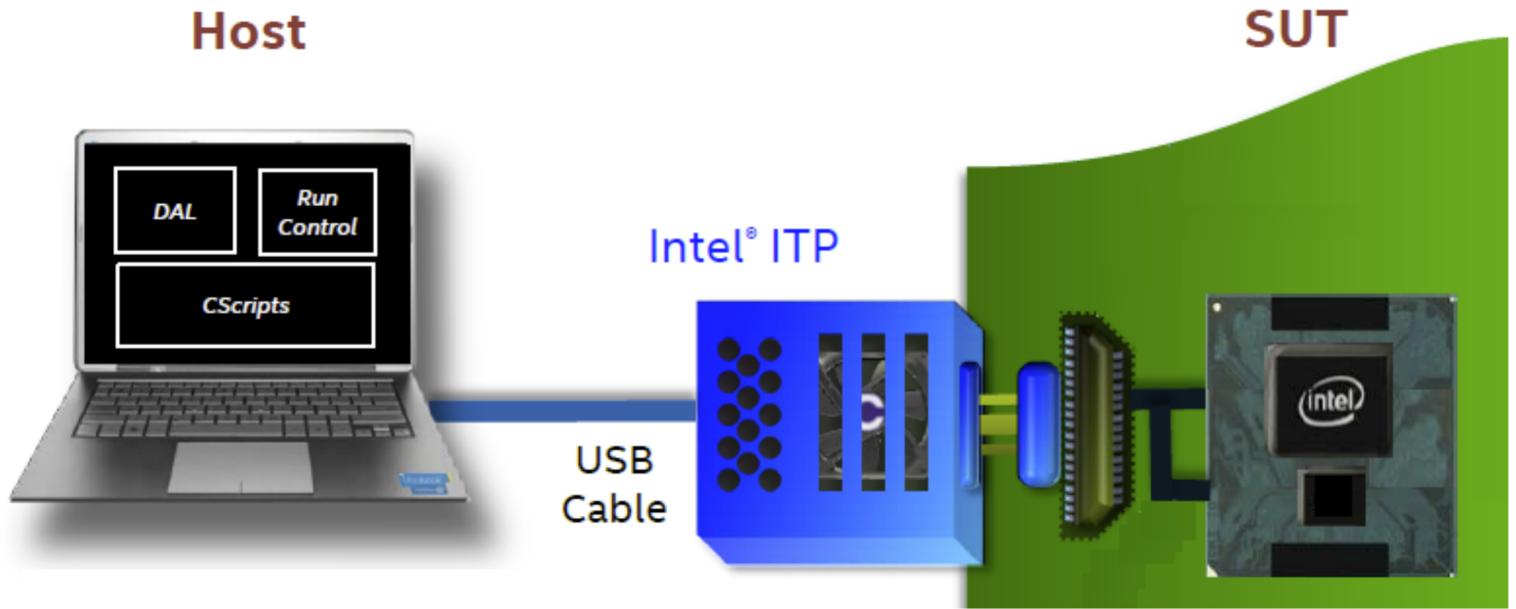
System Information



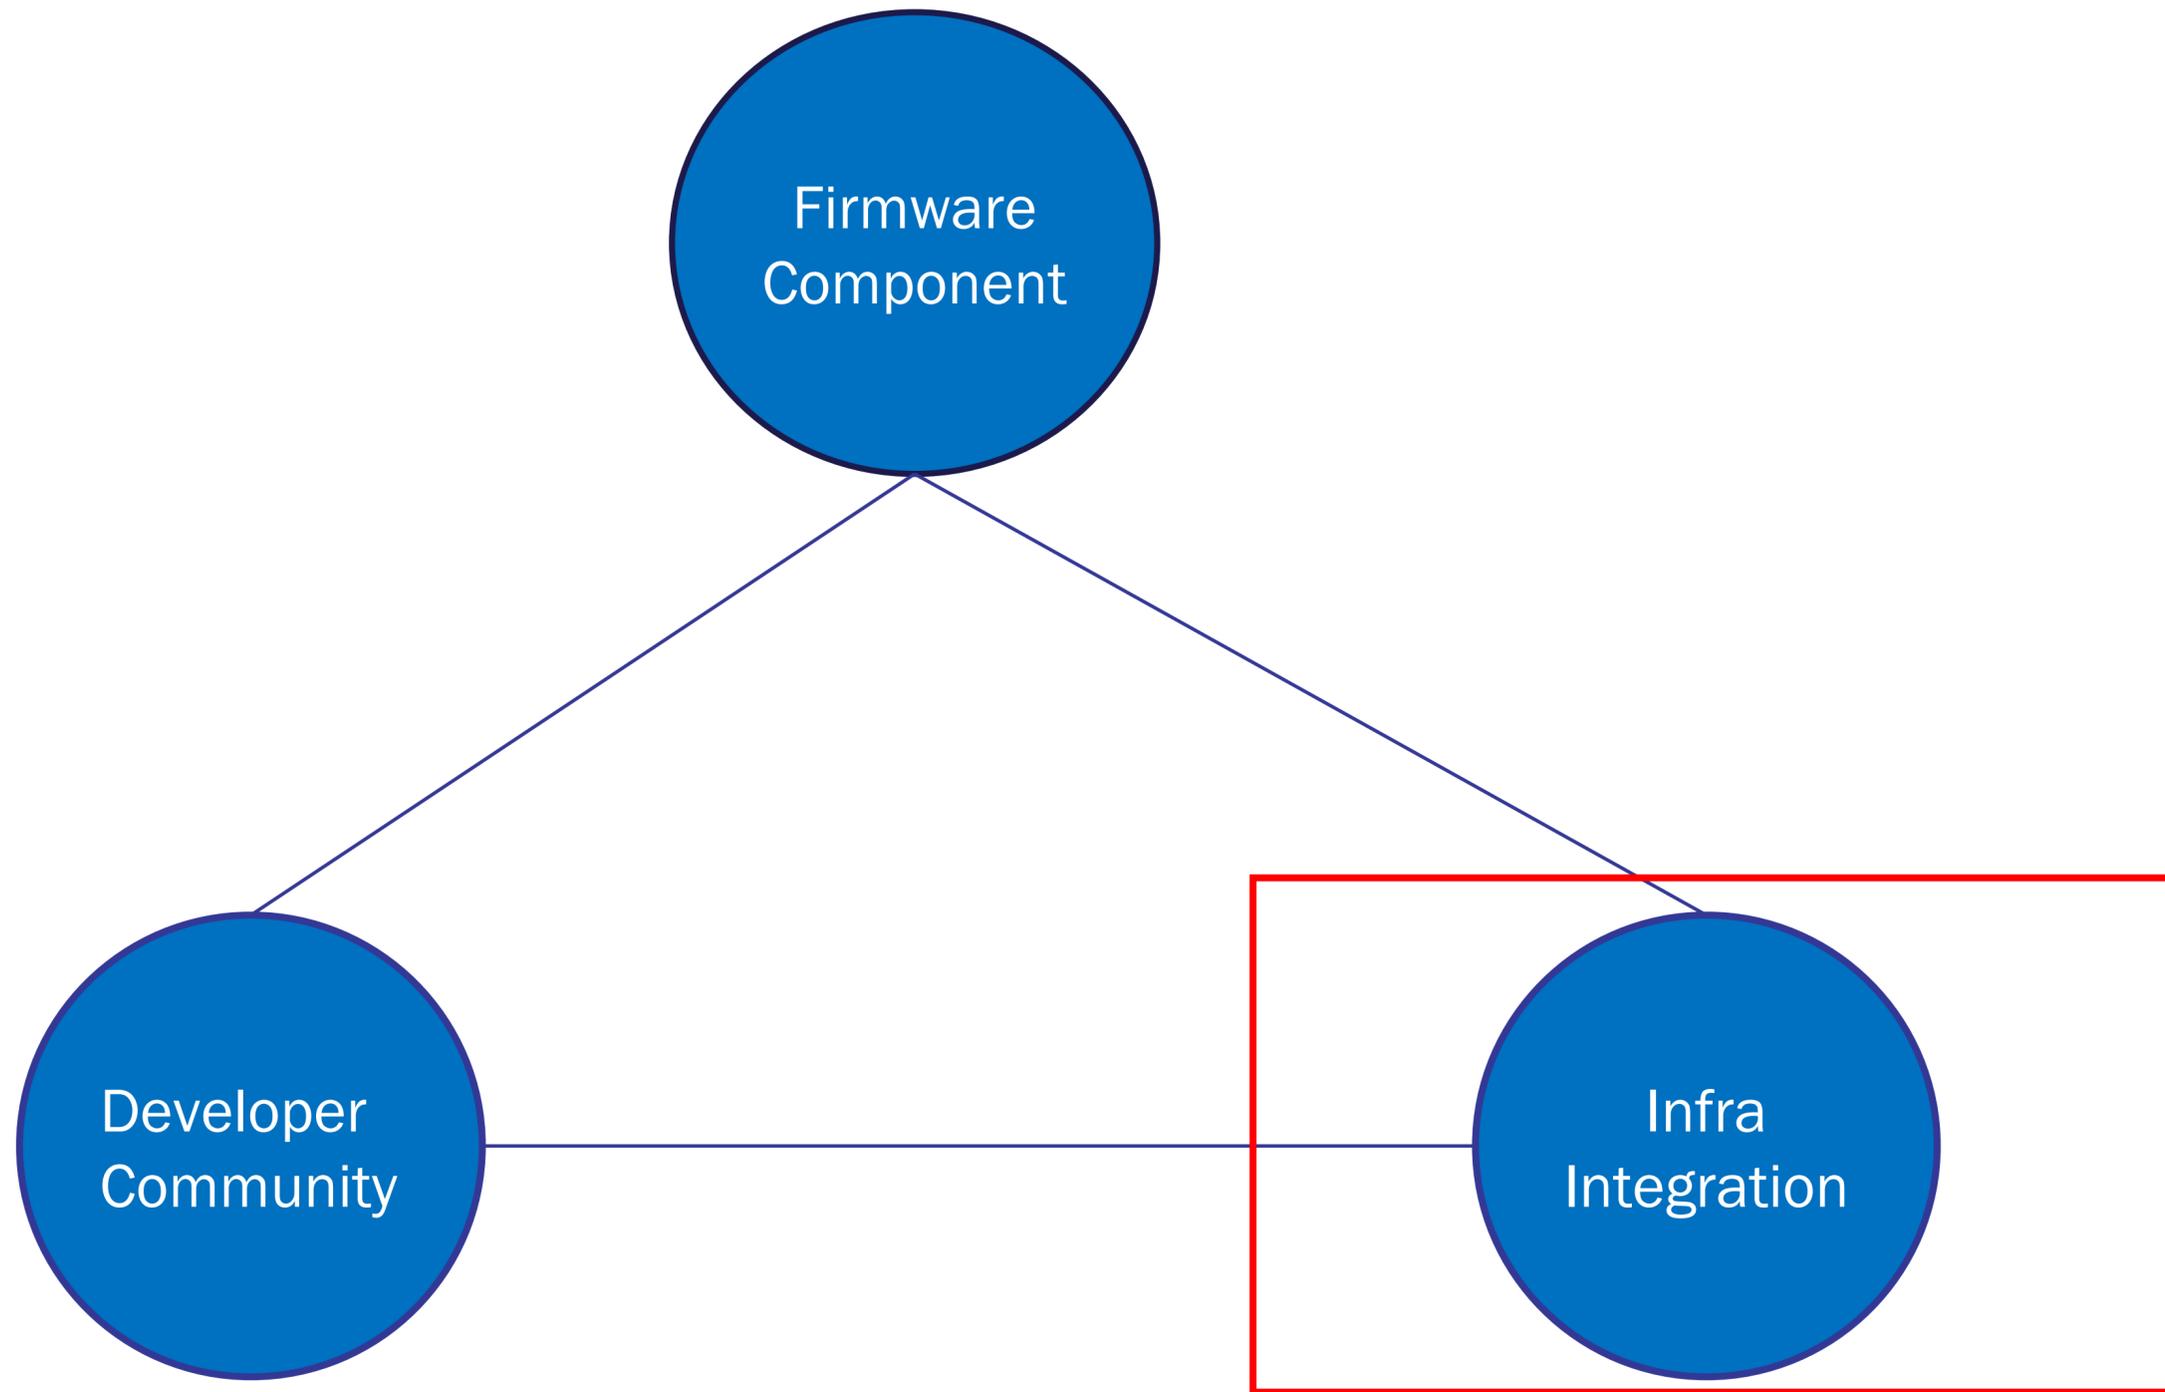
Critical Sensor Info



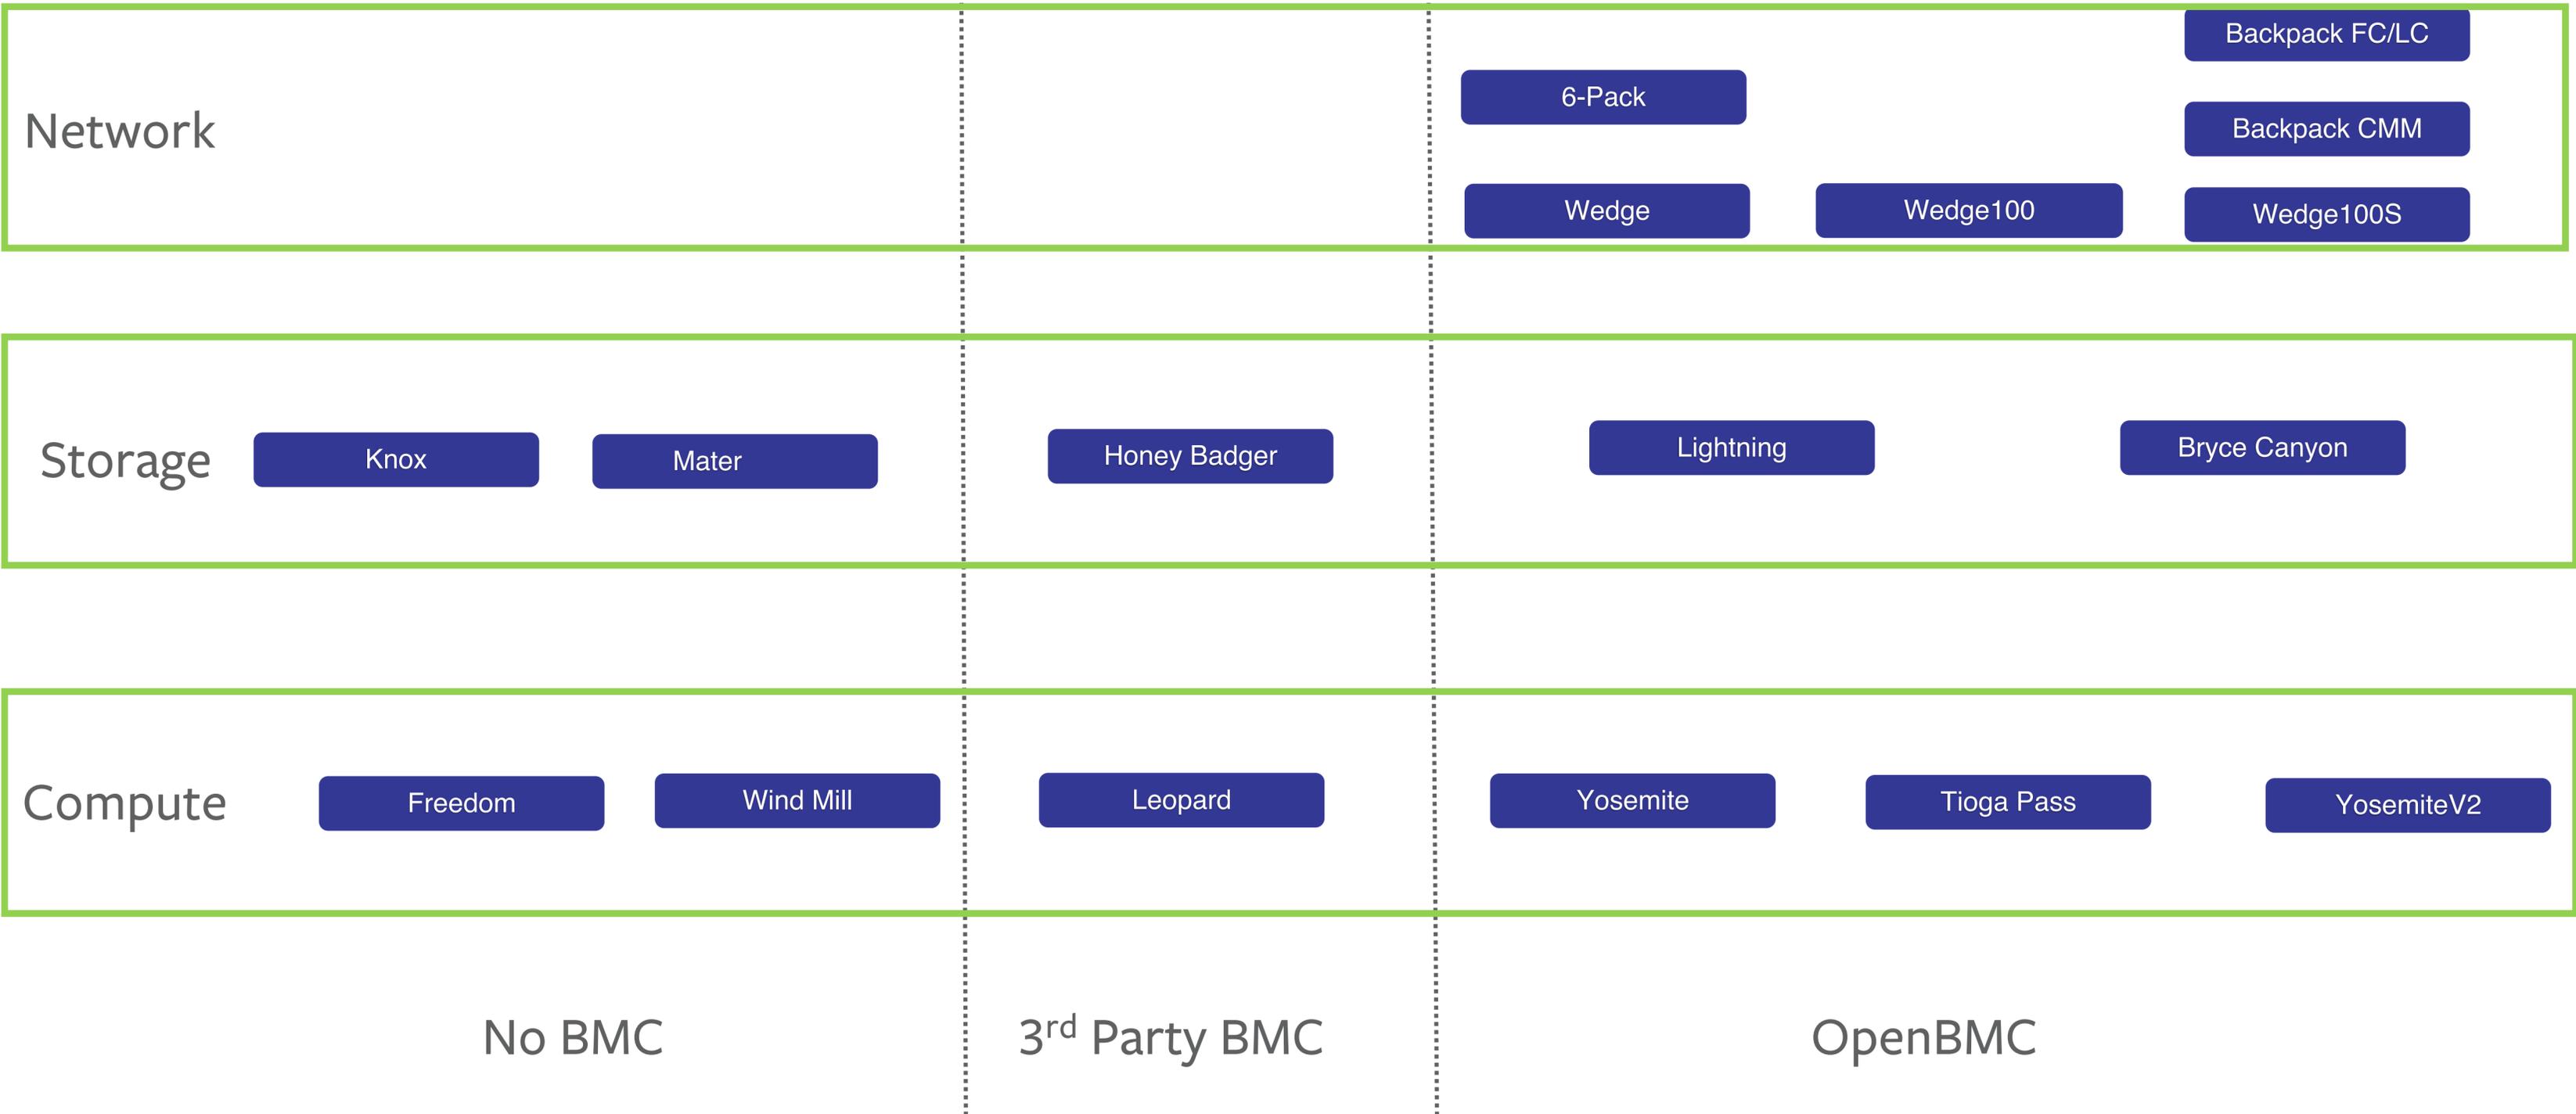
At-Scale Debug



*Image courtesy of ASSET InterTech



OpenBMC@FB



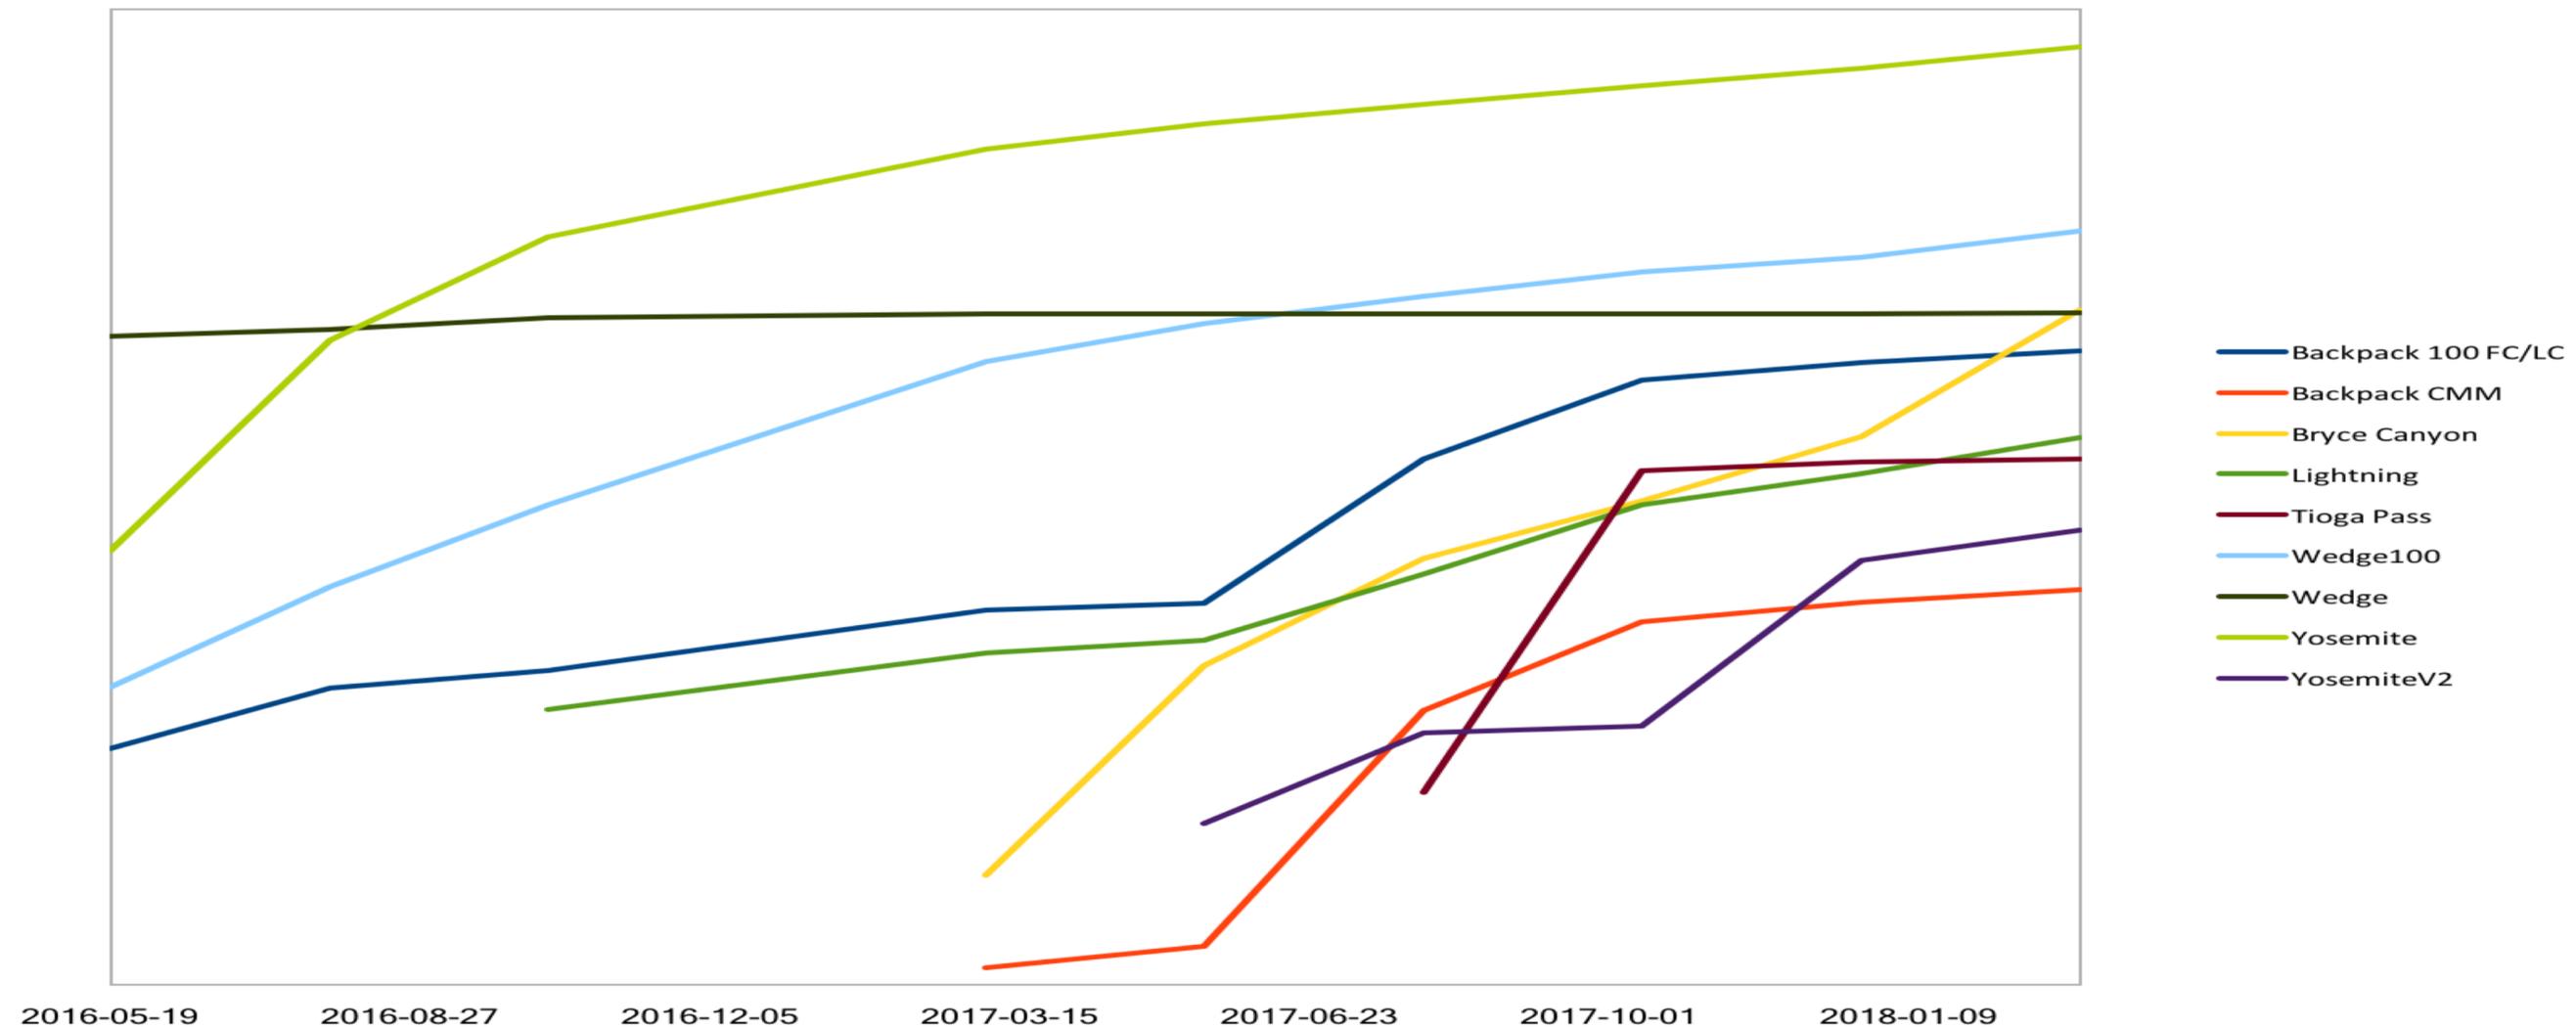
No BMC

3rd Party BMC

OpenBMC

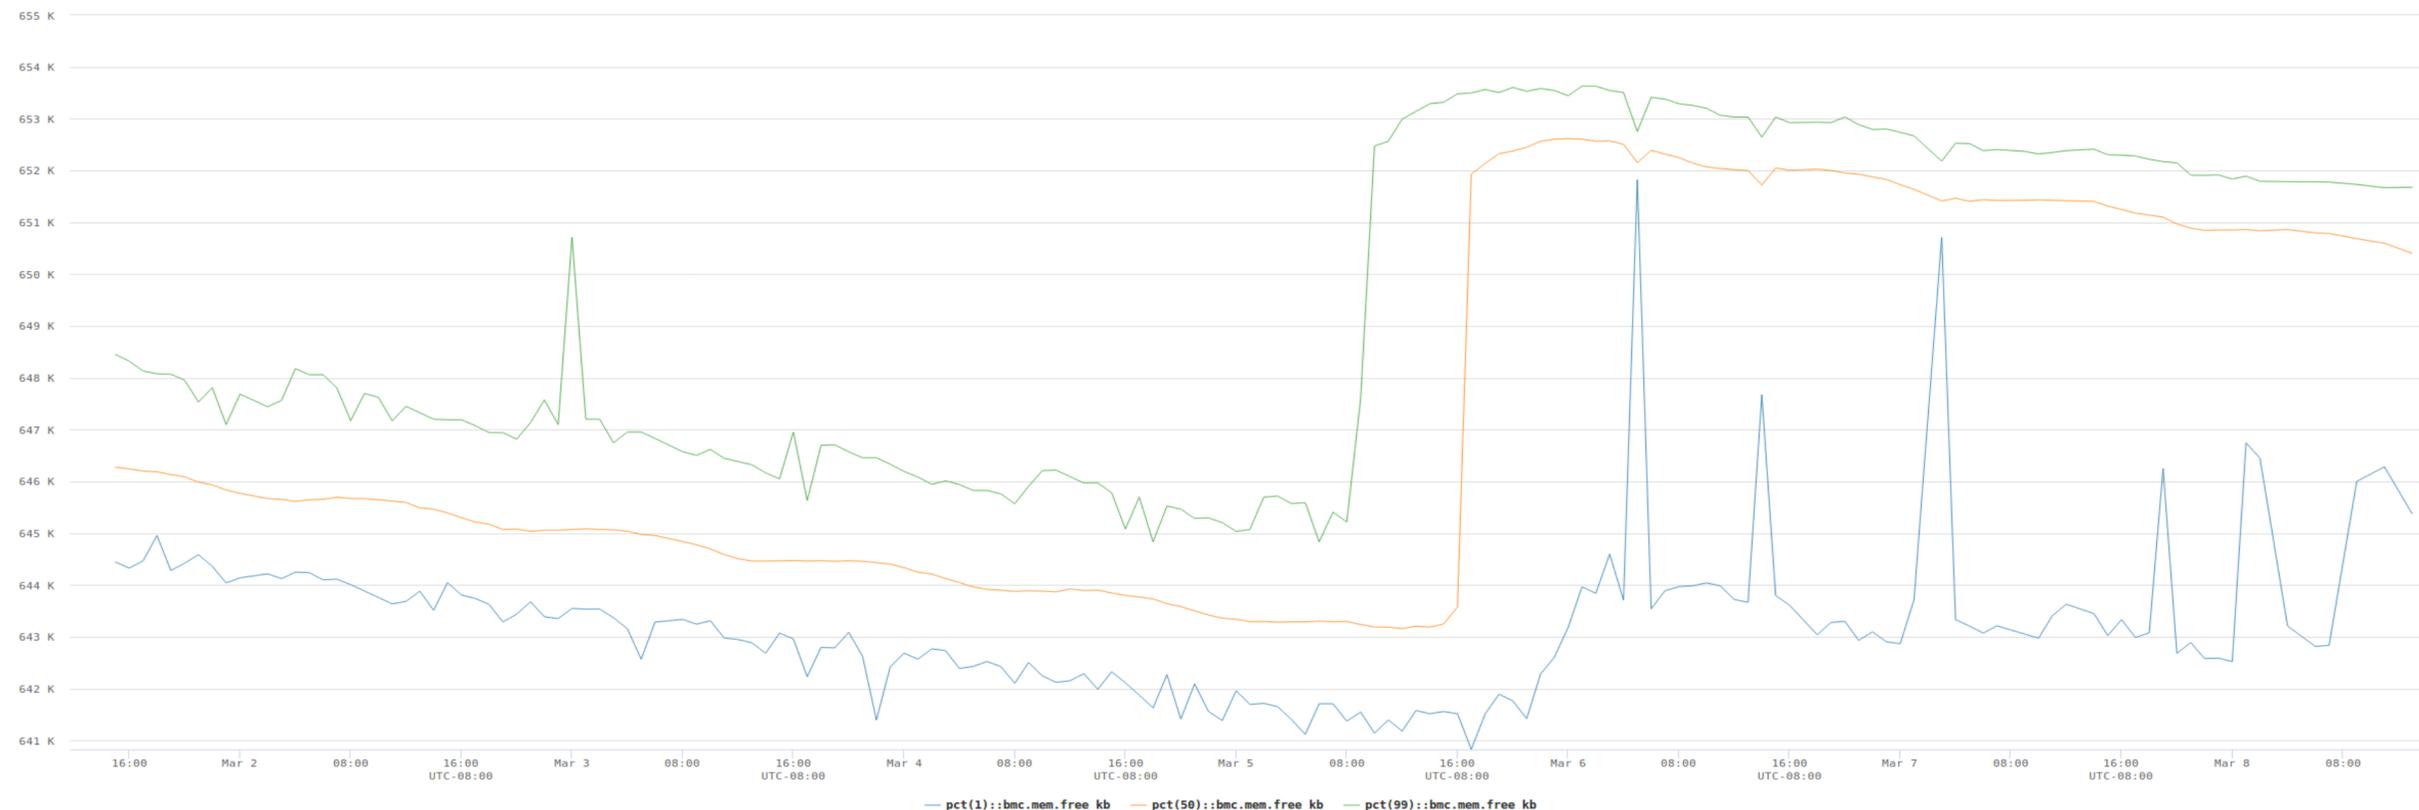
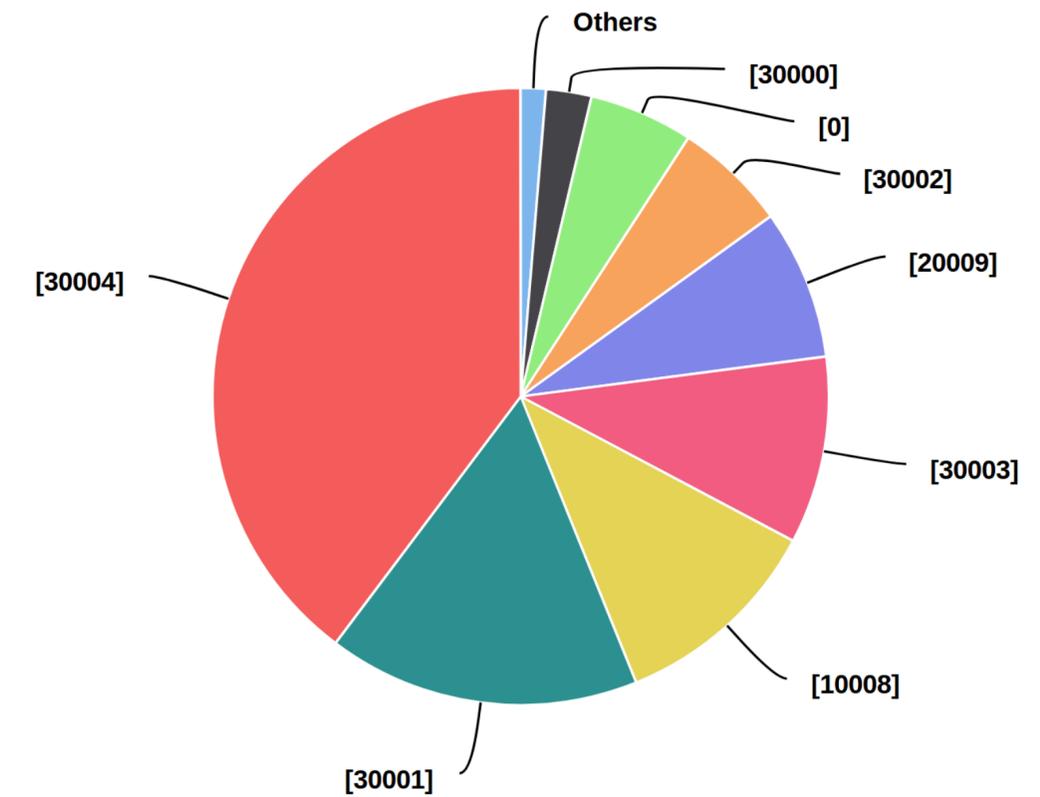
Monitoring, Updates, and Configuration

- Growing number and variety of deployed
- Trending away from host-based/in-band implementations
- Leverage OpenBMC as a capable Linux system



Monitoring OpenBMC

- Periodic fleet-wide polling of REST APIs
 - Beringei time-series database
 - Base image version, memory usage
- Syslog pipelines logs
 - Host serial console
 - OpenBMC system logs



Updating OpenBMC

- Improved resilience and speed
 - Improved unit test coverage
 - More prerequisite checks
 - More nohup
 - Less back-and-forth over SSH
 - Default to HTTP (with checksum validation) instead of SCP
- Queue of out-of-date systems computed from version data from monitoring

```
INFO:2018-02-14 16:38:04,733 132496 KiB free memory.
INFO:2018-02-14 16:38:04,894 Validating partitions in u-boot (/dev/mtd0, 0x60000 bytes), env (/dev/mtd1, 0x20000 bytes), kernel (/dev/mtd2, 0x200000 bytes), rootfs (/dev/mtd3, 0xc00000 bytes), data0 (/dev/mtd4, 0x1100000 bytes).
INFO:2018-02-14 16:38:06,089 0x0060000@0x0000000(u-boot) has known good md5sum.
INFO:2018-02-14 16:38:06,143 data_crc32: 0xaa0307eb
INFO:2018-02-14 16:38:06,645 0x0020000@0x0060000(env) has valid data crc32.
INFO:2018-02-14 16:38:06,816 load_address: 0x40008000, magic: 0x27051956, name: Linux-2.6.28.9, image_type: 2, cpu_architecture: 2, header_crc32: 0x982fa4f5, creation_time: 2017-02-15, compression_type: 0, data_crc32: 0x86689ef6, entry_address: 0x40008000, os: 5, data_size: 1732472
INFO:2018-02-14 16:38:10,117 0x0280000@0x0080000(kernel) has valid data crc32.
INFO:2018-02-14 16:38:10,171 load_address: 0x40800000, magic: 0x27051956, name: yosemite-image-yosemite-20170215, image_type: 3, cpu_architecture: 2, header_crc32: 0x20588132, creation_time: 2017-02-15, compression_type: 3, data_crc32: 0xc0a56c2b, entry_address: 0x40800000, os: 5, data_size: 12155837
INFO:2018-02-14 16:38:27,367 0x0c00000@0x0300000(rootfs) has valid data crc32.
INFO:2018-02-14 16:38:48,519 0x1100000@0x0f00000(data0) readable.
INFO:2018-02-14 16:38:49,155 Kernel parameters before changes: debug console=ttyS0,57600n8 root=/dev/ram rw
INFO:2018-02-14 16:38:51,820 Kernel parameters after changes: debug console=ttyS0,57600n8 root=/dev/ram rw mtdparts=spi0.0:0x0060000@0x0000000(u-boot),0x0020000@0x0060000(env),0x0280000@0x0080000(kernel),0x0c00000@0x0300000(rootfs),0x1100000@0x0f00000(data0),-@0(flash0)
INFO:2018-02-14 16:38:53,392 Proceeding with reboot.
```

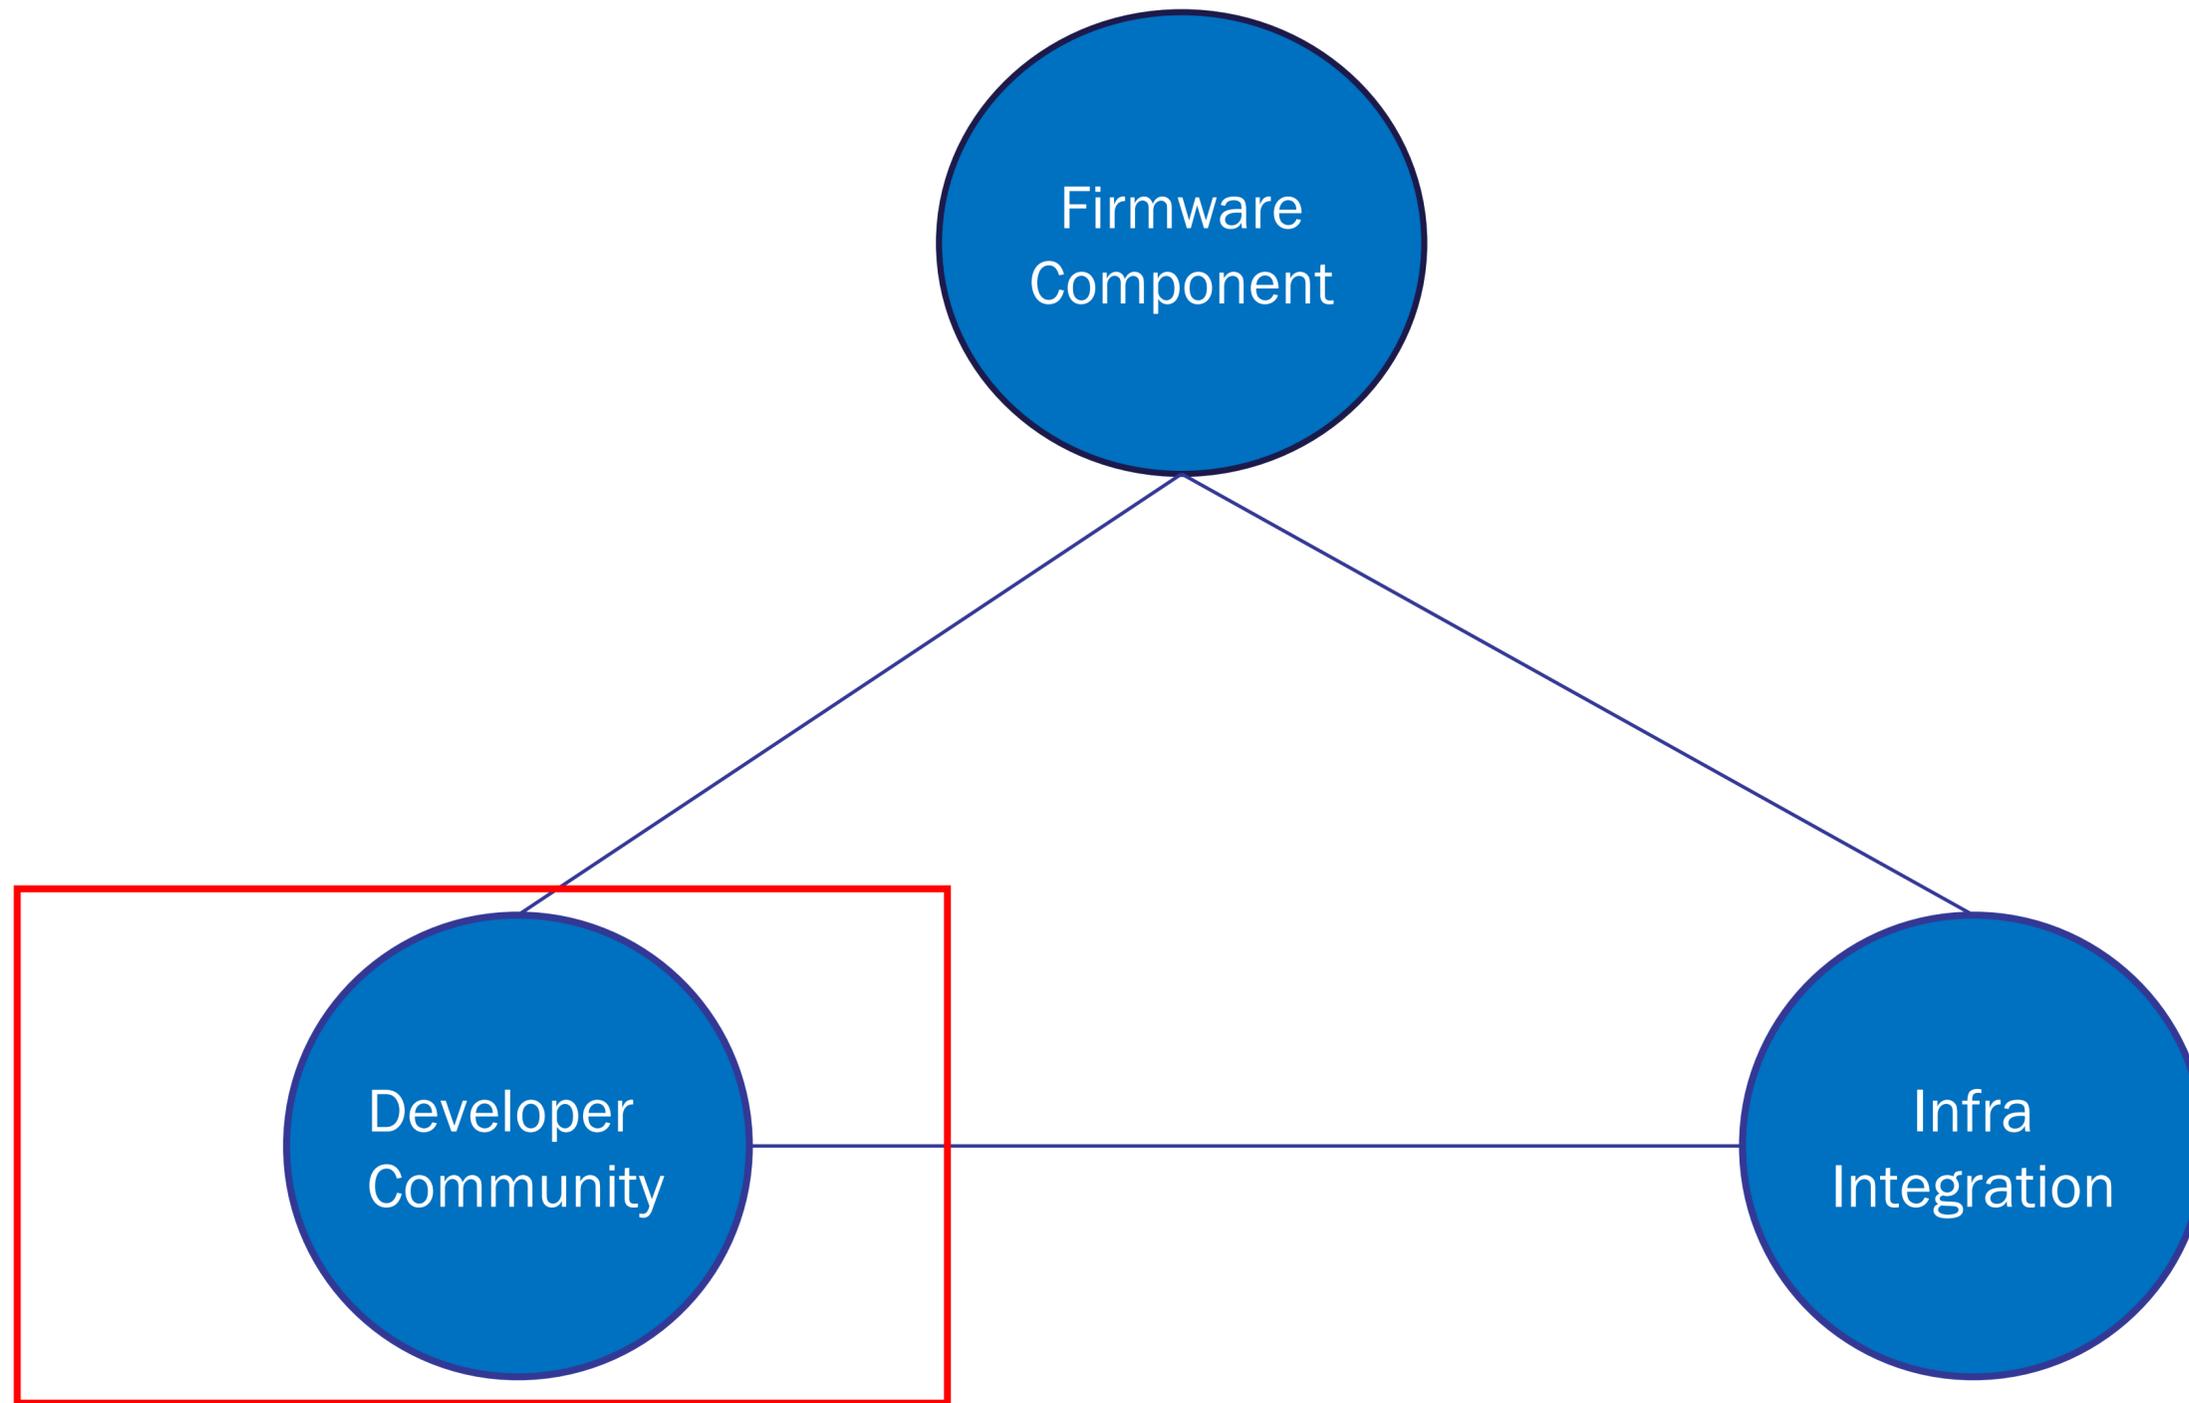
Configuring OpenBMC

- Host-based Chef and simple tarball solutions in use on different portions of fleet
- Experimented with Ansible
- Working default configurations now included in base images at Yocto build time

```
3 meta-facebook/recipes-support/ntp/ntp_%.bbappend View v
... @@ -0,0 +1,3 @@
1 +FILESEXTRAPATHS_prepend := "${THISDIR}/files:"
2 +
3 +SRC_URI += "file://ntp.conf"
```

Configuring OpenBMC

- SLAAC provides initial IPv6 address
- DHCPv6 provides assigned IPv6 address, DNS servers, DNS search path
- NTP and rsyslog use DNS short names



Developer Community

- OpenBMC Distributions
 - <https://github.com/openbmc>
 - <https://github.com/facebook/openbmc>
 - Multiple **private** derivative distros!!
- Goal
 - Shared Development
 - Industry standard Interfaces
 - Unified OpenBMC Distribution
- Community Project
 - Major Partners
 - Linux Foundation Project





OCP SUMMIT