



# OPEN

Compute Project

## Cell Site Gateway Router

Revision 1.5

Authors: Layli Amiri, Tuan Duong, Wen Chien

## Acknowledgement

We would also like to acknowledge our colleagues at AT&T and suppliers and providers who provided invaluable insights into how the cell site gateway router can be used. Lastly, the OCP, without whose support and existence this type of work would not be possible.

## Revision History

Revision	Date	Author	Description
1.0-1.4	9/17/2018	Tuan Duong	Initial Release & Internal Reviews
1.5	9/21/2018	Tuan Duong	Initial Release to OCP

## Contents

Acknowledgement .....	2
Revision History .....	2
Contents .....	2
Licenses .....	4
Scope .....	4
System Overview .....	4
Use Case Topology .....	5
System Functional Block Diagram .....	5
Physical Design Constraints .....	6
Hardware Compliance Requirements .....	6
Power Supply Specifications .....	7
Dying Gasp Guidance .....	7
Fan Module Specifications .....	7
Internal Glue Logic and Controls .....	7
Port Numbering Specifications .....	8
LED Operations Recommendations .....	9
Silk Screen Recommendations .....	10
Interface Specifications .....	10
Timing Circuitry Design Specifications .....	11
MAC Specifications .....	12
Number of MAC addresses and Address Constraints .....	12
Mac Address Conservation .....	12
X86 CPU Specifications .....	12
Trusted Platform Module (TPM) (Optional) .....	12
Coin-Cell Lithium Battery .....	13

Watchdog Implementation .....	13
HW BMC Specifications .....	13
Thermal Shutdown .....	14
Software Support.....	14
ONIE.....	14
BMC Software .....	14
Network OS (NOS).....	14

## Licenses

This specification is contributed under the OCP Contributor Licensing Agreement (OCP-CLA) by AT&T.

Limitations of the OCP CLA license are noted below:

All devices that may be referred to in this specification, or required to manufacture products described in this specification, will be considered referenced only, and no intellectual property rights embodied in or covering such devices shall be licensed as a result of this specification or such references. Notwithstanding anything to the contrary in the OCP-CLA, the licenses set forth therein do not apply to the intellectual property rights included in or related to the devices identified in this specification. For clarity, no patent claim that reads on such semiconductor devices will be considered a "Granted Claim" under the applicable OCP-CLA for this specification.

You can review the signed copies of the OCP-CLA for this specification on the OCP website. <http://www.opencompute.org/products/specsanddesign>

Usage of this specification is governed by the OCPHL permissive.

You can review this license at <http://www.opencompute.org/participate/legal-documents/>

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.

## Scope

This document defines the high level technical specifications for the temperature hardened OSP (Out Side Plant) *Cell Site Gateway Router* (CSGR) that AT&T is submitting to the Open Compute Project. The specifications list key requirements and constraints that the design must meet but leaves room for innovation for the hardware manufacturer in the design/implementation/manufacturing process.

## System Overview

This document describes the technical specifications of the *Cell Site Gateway Router* (CSGR). The system is a physical unit that is self-contained, and not part of a rack or other physical chassis system. The system is typically deployed in an ambient environment specified in Telcordia GR-3108 Out Side Plant (OSP) Class 2. The system MUST meet AT&T TP76200 and TP76450 Level 3 requirements for OSP Class 2 deployments<sup>1</sup>. For reference, Telcordia GR-487 specifies enclosure cabinets that would provide/maintain the GR-3108 Class 2 ambient operating conditions and within which this device would be deployed.

The CSGR is a high-performance, versatile router designed to address the changing needs of the backhaul

---

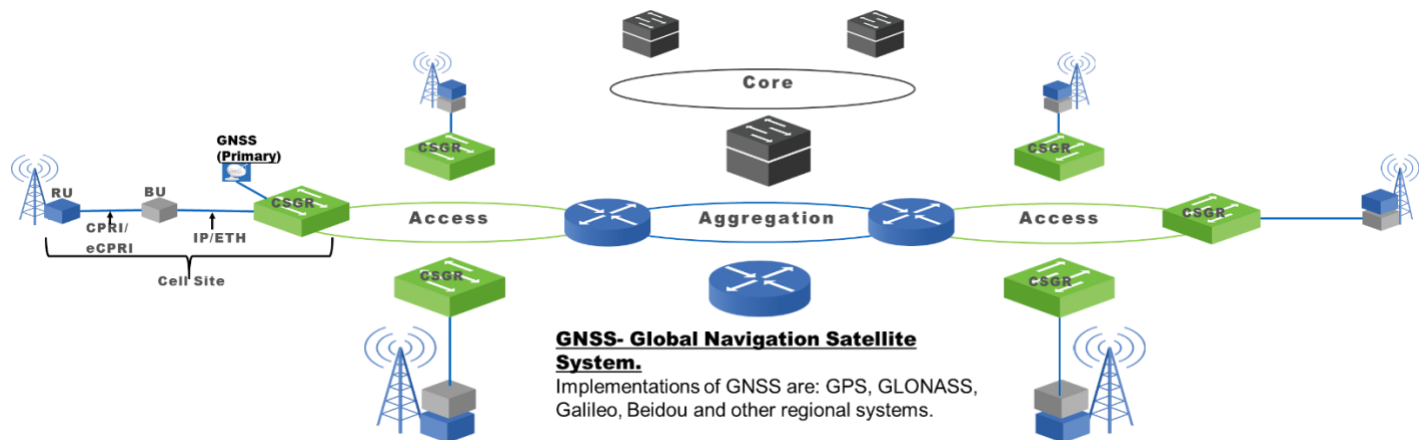
<sup>1</sup> AT&T Technical Publications are available at: <https://ebiznet.sbc.com/sbcnebs/Intropage2.htm>

transport requirements as the Mobile Service Providers start to make the transition from legacy 2G, 3G, 4G toward 5G RAN technologies. It has an open hardware with disaggregated software solution offering the following key advantages:

- It decouples hardware from software, so providers have freedom to choose software implementations that better support their business and operational models.
- It is designed to support {100M/1G} connections to legacy Baseband Unit (BU) systems to next generation 5G BU systems that can operate at {1G/10G/25G} and backhaul speeds transitioning from 1G to {10G/25G/100G}.
- It is designed to operate at Industrial temperature range (-40C to + 65C ambient) meeting AT&T TP76200/TP450 Level 3 requirements for sealed OSP Class 2 cabinets.
- It features the Broadcom Qumran-AX MAC with deep buffers to support advanced QOS features.
- It includes an innovative onboard BMC to enable OOB Management Ethernet to accommodate Zero Touch Provisioning and onboarding models.
- It includes an atypically powerful CPU with 20G of bandwidth to the switch. This can support VNFs in an OS that perform significant signaling, diagnostic or data collection, and even those that perform functions for a fraction of the data plane.
- A separate timing circuitry block supports a variety of Timing Inputs {GNSS, TOD, T1/E1-BITS, 1PPS, 10Mhz} and Timing Outputs to {1PPS, 10Mhz} to adapt to the evolving timing requirements and implementations in the 5G technology evolution.
- TP76200/TP450/NEBS Level3 Independent Lab Certification.

## Use Case Topology

The following diagram shows a high-level overview of the topology where the Open CSGR can be used.



**Figure 1 – Cell Site Gateway Router (CSGR) Use Case**

Since cabinet space, power, cooling is a premium at cell sites, it is envisioned that the initial deployment for the CSGR will be existing Cell Sites that support current technologies (3G,4G) with additional support for newer 5G. Hence, support for a variety and number of interfaces are built into the system.

## System Functional Block Diagram

Figure 2 shows a block diagram of the major sub-system components: CPU, BMC, MAC, Timing, Power, FAN,

Dimensions, Airflow, Interfaces and Interface types.

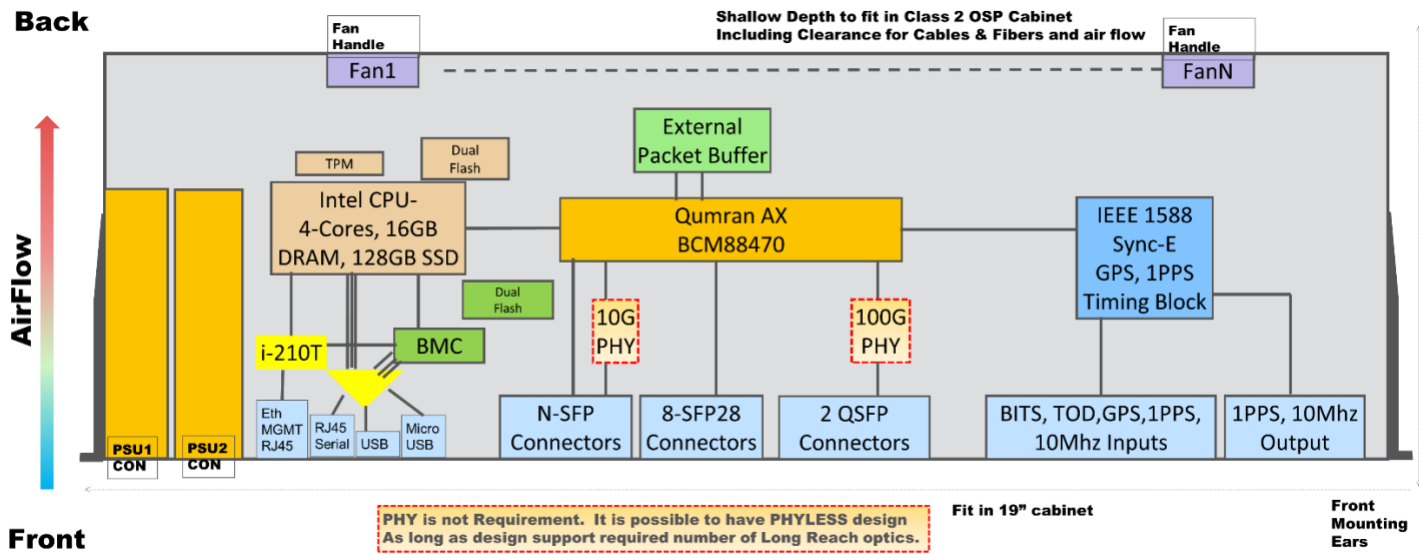


Figure 2 – AT&T Cell Site Gateway Router Block diagram

## Physical Design Constraints

With the proper cabling the following consideration must be made for the physical platform dimension (does not include handles or SFP). All cabling must be front accessible and adhere to AT&T bend radius standards as specified in ATT-TP76300, section J part 2.10.

- **Width:** 19" rack mount EIA cabinet standard (must support both 2 and 4 post mounting)
- **Depth:** 11.9" must fit in 17.9" cabinet including all cabling and air flow clearance
- **Height:** 1RU
- **AirFlow:** Must be front to back.
- **Access:** All Cabling and power cabling must be Front Accessible. (There is no rear access in OSP Class 2 cabinets)
- **Temperature:** -40C to + 65C Ambient. As such, all components selected must operate within this temperature range. Where there are no components available to operate within this range, then mitigation measures must be designed into the system to enable it to operate within this range.

## Hardware Compliance Requirements

The Open Cell Site Gateway Router MUST meet the following requirements:

- AT&T TP76200 (Issue 20) & TP76450 (v17) for Level 3 for an environment specified in GR-3108 OSP Class 2.
- Typically a GR-487 compliant remote cabinet will provide an internal ambient operating environment meeting GR-3108 Class2 specifications.

- Copies of this document and general information about AT&T's environmental equipment standards can be found at <https://ebiznet.sbc.com/sbcnebs/>

## Power Supply Specifications

- The power supplies shall accept nominal -48V DC Power supply.
- The power supplies shall meet the 80 Plus standard for high efficiency - with higher levels desired.
- The system shall have less than 350W total power consumption.
- There shall be redundant and field replaceable power supplies.
- Hot swap not required. However, if design supports hot swap capability, then it would be a plus.
- Power connectors at the front panel with screw-in single hole lug connections.
- The power supplies should support loss of power indicators to facilitate BMC
- Front access grounding post.

## Dying Gasp Guidance

Dying Gasp is not a required feature for this application. However, if dying gasp can be cost-effective designed into the system, then this would be a plus because the system can be potentially used in other use cases where dying gasp is required.

In Dying Gasp, moments before the power is completely lost to the system, the system should send out an alarm message to an upstream management system that it is about to lose all power.

Where this feature is critical are in locations where the service providers has no other means to identify the loss of power. An example would be in a customer premise. In a central office or a service provider controlled environments, there are typically other means to determine if a site has lost power.

A dying gasp implementation needs to have some sort of short term energy storage device and circuitry and PCB traces to maintain power to critical system components so that the NOS have enough time to generate and transmit the alarm. For this to be useful to a service provider application, the message needs to be transmitted to the upstream management system via the Uplink connection to the device (like the 10G or 100G interfaces) rather than the BMC OOB management ethernet interface. How much time is dependent upon NOS implementation. Some rough guidelines are 30-50msec.

## Fan Module Specifications

Redundant and field replaceable (when rear access is available). Hot swap not required but if hot swap is supported then it would be a plus (+). All front panel space will be occupied by the interfaces and PSU. In a OSP Class 2 cabinet, there is no rear door access. So to replace a fan in this situation would require a dismounting of the unit. However, when used in other applications that has rear door access, then FANs can be field replaceable.

## Internal Glue Logic and Controls

The specification leaves freedom to manufacturer to use any combination of discreet logic/PLD/CPLD/FPGA necessary to implement of the specification. It is recommended to use current practice and provide I2C

interface to the Host for control of PSU, FAN, Optics, and any other key components such as reading of registers or erasing and flashing of NVRAM, EEPROM, Flash,..... The manufacturer must provide the instructions and drivers to access these components as part of the BSP- Baseboard Support Package, so that NOS development can access and control these components as necessary.

## Port Numbering Specifications

AT&T numbering standard for white boxes starts from zero {0,1,2,...}, increasing from Left to Right. This applies to Ports, FAN and PSU and other FRU (Field Replaceable Units). Numbering starting from Zero also applies to break out ports which is more dependent upon software implementation as opposed to hardware and silk screening implementations.

Manufacturer has a degree of freedom for the numbering with respect to vertical grouping. For examples, the following schemes are acceptable.

**NOTE:** The port numbering illustration shown in these tables does not reflect the actual number ports specified for the Open CSGR. It is just to illustrate the acceptable numbering scheme.

0	2	4	6	8	10	12	14	16	18
1	3	5	7	9	11	13	15	17	19
20	22	24	26	28	30	32	34	36	38
21	23	25	27	29	31	33	35	37	39

**Table 1:** 2 grouping numbering scheme: Upper/lower port grouping. Sequentially Numbered Upper grouping then followed by Lower grouping.

0	1	2	3	4	5	6	7	8	9
10	11	12	13	14	15	16	17	18	19
20	21	22	23	24	25	26	27	28	29
30	31	32	33	34	35	36	37	38	39

**Table 2:** 4 rows numbering scheme: Sequentially number each row then move to next row.

0	4	8	12	16	20	24	28	32	36
1	5	9	13	17	21	25	29	33	37
2	6	10	14	18	22	26	30	34	38
3	7	11	15	19	23	27	31	35	39

**Table 3:** 1 Grouping numbering scheme. Sequentially number the ports in a column then move to next column.

As shown in the system block diagram, the numbering for the SFP groups of ports should start from 0 and groups of QSFP form factor ports should also start from 0. The table below illustrates this concept.



0	2	4	6	8	10	12	14	16	18		0	2	4
1	3	5	7	9	11	13	15	17	19		1	3	6

**Table 4:** 0-19 are ports of one physical grouping characteristics (e.g. SFP Pluggables). 0-5 are ports of different physical grouping characteristics (Example: QSFP-Pluggables)

## LED Operations Recommendations

Refer to AT&T Hardware Common Systems Requirements for recommendations on system and interface LED colors and operations.

The indicator lamps (LEDs) must convey the information described in 4. The number, colors, and flash behaviors are desired but not mandatory.

**Table 4 - LED Definitions (Recommendations)**

LED Name	Description	State
PSU1	Led to indicate status of Power Supply 1	Green - Normal Amber - Fault Off – No Power
PSU2	Led to indicate status of Power Supply 1	Green - Normal Amber - Fault Off – No Power
System	LED to indicate system diagnostic test results	Green – Normal Amber – Fault detected
FAN	LED to indicate the status of the system fans	Green – All fans operational Amber – One or more fan fault
LOC	LED to indicate Location of switch in Data Center	Blue Flashing – Set by management to locate switch Slow flashing – System is in standby state Off – Function not active
SFP- LEDS	LED built into SFP(28) cage to indicate port status	On /Flashing – Port up (flashing indicates activity) Green – Highest Supported Speed Amber – Lower Supported Speed Off – No Link/Port down
QSFP LEDs & Breakouts	Each QSFP28 has four LEDs to indicate status of the individual 10-25G ports	On Green/Flashing – Individual 25G port has link at 25G. (yellow for 10G) Green – Highest Supported Speed Amber – Lower Supported Speed Off – No Link/Port down
OOB LED	LED to indicate link status of 10/100/1000 RJ45 management port	On /Flashing – Port up (flashing indicates activity) Green – Highest Supported Speed (1G) Amber – Lower Supported Speed (100M/10M) Off – No Link/Port down

Timing LED	??	??
Remote Management LED	LED to indicate that the system has connectivity to central management	On Green – Management System connected Off - no connection

## Silk Screen Recommendations

It is a strong recommendation that the manufacturer choose Pantones, Contrast, and Font Size that MAXIMIZE visibility to the field technicians working in low light, tight spaces and crowded cabling conditions.

Silk screen should be clear and avoid possible confusion or misinterpretations.

Best is to solicit feedback prior to implementation of silk screen.

## Interface Specifications

- **Craft/Management Interfaces**

Only one Serial input can be active for the Console. Micro USB will have higher priority than RJ45 and USB serial by default.

The Serial console needs to support default selectable between the BMC or the X86 CPU. (customer-provisioned choice).

The RJ-45 OOB Ethernet management port needs to be operational even when the system is in the shutdown mode. As such it needs to be designed using the standby power rail. It also needs to provide simultaneous connectivity to the X86 CPU and the BMC.

The Intel I-210 NIC is specified to support DPDK.

<b>Craft Type</b>	<b>QTY</b>	<b>Purpose</b>
Micro USB Serial	1	Console
RJ-45 USB Serial	1	Console
USB Port	1	USB data access
RJ-45 10M/100M/1G	1	Ethernet OOB Powered on Standby Power rail

- **Pluggable Interfaces, Speeds, Reach**

BCM88470 platform supports variety of SERDES that can programmed to operate at different modes and speeds to support different type of network interfaces. The following are design requirements for the front panel interface ports which can serve as a guideline for PCB traces of the PM10Q vs PM25 Serdes on the BCM88470. It gives manufacturer degrees of freedom for PHY versus PHYLESS designs.

<b>Description</b>	<b>MIN QTY</b>	<b>Distance</b>
Number of SFP Pluggable Front Panel Ports	24	NA
Number of 1G/10G SFP pluggable that support Long Reach	8	ZR (up to 80km)

Number of SFP pluggable supporting 1G/10G/25G speeds	8	SR, LR (up to 10km)
Max number of SFP Pluggable interfaces on front panel	N where N>=24	NA
Number of SFP Pluggable Interfaces that support 100M/1G/10G <a href="#">Auto Negotiation capability</a>	N – (minus) 8	NA
Minimum number of QSFP Pluggable Interfaces supporting 40G/100G	2	ER4 (up to 40km)

The key to understanding this section is the [Minimum total number of front panel interfaces](#) supporting the required features, distance, and type.

- **MACSEC Support: (Optional)**

MACSEC is an optional feature for the Open CSGR specification. If manufacturer’s design uses PHY chips, some PHY Chips include MACSEC capability. The following provides some guidelines regarding MACSEC usage:

- The MACSEC capability should be designed into Long Reach Pluggables: 1G/10G (80km) **and/or** 40G/100G (40km).
- The MACSEC capability should support DOT1Q in the Clear in order to operate over different Ethernet Provider service offerings.

- **Supported Optics**

The CSGR MUST not inhibit use of compatible optics.

Refer to AT&T 3<sup>rd</sup> Party Optics List for the specific application/use case.

- **Timing Input/Output Interfaces**

<b>Timing Description</b>	<b>Input/Output</b>	<b>Connector</b>
1pps	Input	SMA
10MHz	Input	SMA
GNSS	Input	SMA
Time of day (TOD)	Input	RJ-45
Building-Integrated Timing System (BITS)	Input	T1/E1 RJ-48
1pps	Output	SMA
10MHz	Output	SMA

## Timing Circuitry Design Specifications

The hardware needs to support the following timing implementation requirements.

- GPS receiver with automatic and manual delay compensation for delays caused by cable, splitter and antenna.
- SYNC-E- can be activated on any Ethernet port.
- IEEE 1588v2 transparent and boundary clock function with design to utilize Qumran-AX full functionality Including: ITU profiles 8265.1, 8275.1 T-BC/GM, 8275.2 T-BC/GM/TSC-A.
- Stratum 3E Internal Oscillator.

## MAC Specifications

- This platform will utilize the Broadcom BCM88470 integrated packet processor and traffic manager single chip switch. High level functions include:
  - - 300Gb/s full duplex integrated traffic manager and packet processor.
  - - 32K programmable wire-rate queues
  - - Deep buffering 24Gb DRAM
  - - Hierarchical QOS with Ingress shaping and congestion management.
  - - IPSEC Engine
  - - IEEE 1588 transparent and boundary clock

## Number of MAC addresses and Address Constraints.

- For the Cell Site Gateway Router, the anticipated number of MAC addresses that needs to be programmed into the system to be made available to the NOS for use with the BCM 88470 is **256**. This number of mac addresses would leave plenty of room for different NOS implementations in a variety of use cases.
- The BCM88470 has specific constraints with respect to the MAC address. The Block of 256 contiguous MAC address allocated to the CSGR must not cross the bit boundary restriction:
  - The first 36 bits of the MAC address must be the same.

## Mac Address Conservation

Currently, there is no effective way to recycle the mac addresses once they are allocated to a system. The primary reason for this is due to the lack of a closed loop in the life-cycle management of the hardware. But in the new model where the providers obtain the equipment from the manufacturer, put them into service, and take them out of service to the destruction phase, there is an opportunity to be responsible and recycle the Mac address usage. This would require some discipline built into the hardware life-cycle management to achieve this objective.

## X86 CPU Specifications

The following are minimum level specifications

- Intel x86 platform Xeon-D1500 series with 4 Core minimum. 1.4 Ghz or better.
- 16GB EEC DDR4 and 128GB SSD.
- Dual flash operating in Primary/Backup mode for recoverable remote field upgrade.
- 2x10GE bus design between CPU and the BCM88470 platform to support high BW packet processing by CPU when needed.
- All NIC must be supported by DPDK (refer to <http://dpdk.org/doc/nics>)

## Trusted Platform Module (TPM) (Optional)

The latest Trusted Platform Module (2.0 or greater) shall be used for secure storage of keys and certificates in a hardware chip and is an integral part of creating a Secure Boot environment so that the device cannot be easily taken over, such as by booting from a USB drive.

If the manufacturer can make a design where this is a factory orderable option, that would be a plus.

This is a future proof design specification because it is dependent on the availability of ONIE secure boot. Initial software releases will operate with TPM disabled in BIOS and use regular ONIE Boot process.

## Coin-Cell Lithium Battery

Typical X86 design specifies the use of a Coin-Cell Lithium battery to maintain the RTC. However, the presence of this battery conflicts with the TP76200/TP76450 requirements. As such, the Cell Site Gateway Router design does not have a Coin-Cell battery.

However, there is an issue with correct TPM operation, should TPM be activated in the future, without the presence of the battery to maintain the RTC following a power loss. A ticket was documented with Intel and a work around in the BIOS is needed from the manufacturer. This is documented in Intel Ticket # [\(00260505\)](#).

Coin-Cell battery is a factory orderable item. As such, other providers may choose to use the battery.

## Watchdog Implementation

For the Open CSGR design, it is crucial that some watchdog implementation(s) be designed into the system. The reason is dispatch is costly. If a system hangs, then there needs to be a mechanism to get the system out of this “hung” state. There are a number of implementations, examples are given below:

1. Dedicated watchdog circuitry.
2. Intel TCO watchdog
3. Or some kind of watchdog implementation between BMC and X86 host.

The choice of implementation is left to the manufacturer as long as this is documented and appropriate drivers or mechanisms to leverage this capability from the NOS is provided.

## HW BMC Specifications

The system will be designed with Baseboard Management Controller (BMC) to allow for remote lights out operations, management and access.

The most important requirement for the BMC is that it must be secure. The BMC will be connected to WAN and/or Internet connections. As shown in the System Block Diagram, the BMC has an Ethernet connection which is a shared external connection to RJ45 OOB Management Ethernet Port with the X86 host.

The goal is to enhance OpenBMC over time to support the required features needed for the operational/business mode in a secure way. Before then the firmware must be capable of disabling this Ethernet access when needed.

The following requirements are specified for the BMC.

- Dual Flash memory to support remote reliable in-band BMC Software upgrade.
- System, Qumran, and Host CPU power management
- Temperature monitoring
- Voltage monitoring

- Fan control
- Reset control
- Programming FPGA/CPLD/and other various flash/BIOS
- Read the Rx loss and other signals from the SFP and QSFP ports
- Host CPU boot up status
- System Identifier, including ability to set user-defined identifier, as well as control of locator lamp.
- Serial number / unique identifier
- Board revision ID
- I2C interfaces to Host CPU, USB, temperature sensors, and voltage controllers.
- Monitoring detect signals – including loss of power from the power supplies.
- Must support IPMI 2.0 host mode to provide the following capability via the IPMI interface:
  - ✓ temperature reading and alarms at 3 levels (minor, major, critical) for Processor modules, Chassis, power supply, fans, Broadcom chipset.
  - ✓ status information for fans, power supply, interface modules, processor modules, fan tray

## Thermal Shutdown

NEBS/TP76200 compliant equipment should have the ability to be configured to shut down when the thermal threshold is exceeded or continue to operate until the equipment fails completely. Configurable means that the “user” can select the thermal overload behavior. This must be set through software. The default should always be to implement equipment shutdown in a thermal event

Shutdown means that all non-essential functions of the chassis are powered off and only temperature monitoring capability remain such that, if the thermal event ends, the chassis will autonomously reboot and restore service. One way of accomplishing this is to have the management hardware command the power supplies to shut off their main outputs but maintain an auxiliary power bus that powers the management/monitoring functionality.

## Software Support

The Open CSGR supports a base software package composed of the following components:

### ONIE

Fulfillment of the ONIE hardware specification as laid out here:

[https://opencomputeproject.github.io/onie/design-spec/hw\\_requirements.html](https://opencomputeproject.github.io/onie/design-spec/hw_requirements.html)

### BMC Software

Open BMC with Redfish implementation is the target platform. Commercial BMC with IPMI 2.0 is acceptable to meet near term needs.

### Network OS (NOS)

The AT&T Open Cell Site Gateway Router is an open hardware platform that any NOS vendor can develop to. Equipment should be available with Open Network Linux, and AT&T expects to develop open source support for this platform over time.

