



OPEN

Compute Project

AT&T Open CPE Network Compute Platform Specification

Revision 2.1

Authors: Dustin Grant, John Gibbons, Tom Anschutz

Revision History

Revision	Date	Author	Description
1.0	4/26/2017	Dustin Grant	Initial Release
2.0	8/30/2017	Dustin Grant	Major update changing focus from exploratory questions/feedback to a more definitive specification of the range of devices proposed.
2.1	1/4/2018	Dustin Grant	Removed codename references

Contents

Revision History	2
Contents	2
Licenses	4
Scope	4
Overview	4
Function Component Selection Considerations.....	5
Compute Function	5
CPU	5
Co-Processors.....	6
Memory.....	6
Storage.....	6
Network	6
BMC.....	7
Trusted Platform Module.....	7
Serial Console	7
Ethernet Switch Function	7
Switch chipset	7
RJ45 Ports.....	8
SFP/SFP+ Ports.....	8
Power over Ethernet.....	8
Internal Expansion Module.....	8
Redundant Power Supplies.....	8
System Specifications	8
Common Requirements	9
Hardware Sizes	10

Micro.....	10
Extra Small.....	11
Small.....	12
Medium.....	13
Large.....	14
Overview of Sizes.....	15
Software Support.....	15
BMC support.....	15
ONIE.....	15
Switch Abstraction Interface (SAI).....	15
Open Network Linux.....	16
Specification Requirements.....	16
Safety.....	16
Electromagnetic Compatibility.....	16
ROHS.....	16

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 technical specifications for the AT&T Open CPE Network Compute Platform submitted to the Open Compute Project.

Overview

This document describes the technical specifications of the AT&T Open CPE Network Compute Platform. The system is a self-contained unit that resembles a typical 1RU x86 server or switch. The system is typically deployed at a customer premise site such as an office building or datacenter and is typically installed in a rack or wiring closet.

The purpose of the system is to be able to run Virtual Machines (VM's) containing multiple Virtual Network Functions (VNF's) such as a Router, Firewall, WAN Accelerator, etc. The system would typically be placed at the edge of the customers network to connect to the WAN and/or Internet.

The system described in this document is a hybrid of a typical x86 based server and a high performance Ethernet switch. The x86 part of the system is used to run the VNF's, and the Ethernet switch part of the system is to provide local wire rate L2 or L3 forwarding.

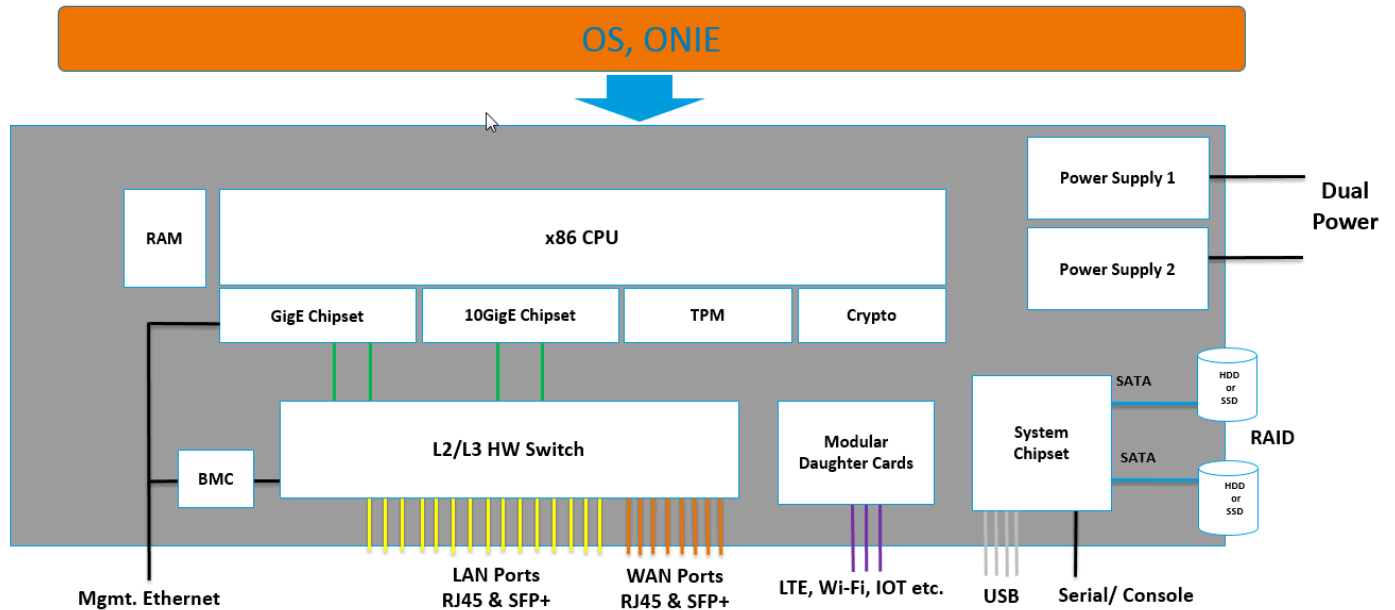


Figure 1 – System Block diagram

Function Component Selection Considerations

This section of the document presents the different functional components, their purpose, and the considerations taken into account when deciding which option to use.

Compute Function

CPU

The CPU is an x86 based CPU that is ultimately responsible for running all of the software functions on the box, including the Base Operating System, Hypervisor, VNF's and also the Control Plane for the Ethernet Switch.

Considerations for choice of CPU:

- Selection of a processor family that can scale from small to large CPE performance requirements (2 cores – 24 cores)
- Single socket vs Dual socket – performance impacts of using a dual socket system vs potential cost savings (ie, 2 6 core CPU's is less expensive than a single 12 core CPU, but there is a performance cost)
- Virtualization Features (VT-x, VT-d, Hyperthreading, etc)
- SoC vs traditional CPU
- Performance
- Cost

For this specification, we are recommending 2 families of Intel x86 processors: Intel Atom C3000 and Intel Xeon-D.

The Intel Atom family is targeted for lower-end platforms and Xeon-D for higher end platforms.

Co-Processors

Support for co-processors for Encryption allows for offloading encryption from the main CPU, freeing the CPU to perform the higher value functions.

The Atom and Xeon-D CPU families both have Intel Quickassist (QAT) built-in to the SoC.

Memory

Memory is relatively straight forward. The memory architecture & type is largely defined by the choice of CPU. The amount of memory should be flexible and sized per the requirements of customer. We recommend always using ECC memory due to the 24x7x365 nature of CPE devices and location in environments with higher levels of interference.

Storage

The amount and type of storage can vary widely based on the use case and size of the customer site. Storage should be modular and customizable per the customers requirements.

We recommend the use of either redundant drives, or the use of recovery drive for cost optimized platforms. The recovery drive is a separate storage device that contains a backup OS to assist in recovery operations or for temporary restoration of services.

Network

There are several types of Network Interfaces, both internal and external. Internally, the x86 Server part of the system will connect to the Ethernet Switch part of the system using Ethernet interfaces. Externally, there will be several types of interfaces: WAN interfaces, LAN interfaces, Management interfaces.

The x86 part of the system should have the following interfaces:

- 2x 10Gbps internally connected to the Ethernet switch
- 2x 1Gbps internally connected to the Ethernet switch
- 1x 1Gbps internally connected to the BMC

The Ethernet Switch ports will be discussed later in this document.

BMC

The system should have a Baseboard Management Controller (BMC) to allow for remote lights out management and console 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 2 Ethernet connections:

- Shared external connection to Management Ethernet Port. This connection is the primary method for connecting to the BMC, and functions much the same way as a BMC does today with a shared Ethernet Port between the BMC and the Host.
- Internal Connection to the Ethernet Switch. This connection is an additional connection to the Ethernet Switch. The purpose of this connection is to allow for multiple Network Paths to the BMC to allow for redundancy. For example, the primary path to the BMC may be over a broadband Internet connection, and a backup path may be over an wireless LTE connection.

Trusted Platform Module

The Trusted Platform Module is 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.

Serial Console

The serial console is traditionally found on network equipment and serves as a last resort for troubleshooting or recovery. With the presence of a BMC, the need for a serial console can be debated. We believe the Serial console still serves a purpose with the presence of a BMC, most importantly when the BMC itself needs to be recovered (password recovery, changing IP address, etc), but it can also compliment the BMC in environments where a traditional serial console server is used for remote console access, or for field technicians to gain local access.

Ethernet Switch Function

Switch chipset

The choice of the Ethernet switching chipset should have the following considerations:

- Number of ports
- Speed of ports
- Range of product family (small to large)
- Programming API (Open vs Closed)
- Feature set (L2, L3, L4+ capabilities)

RJ45 Ports

The bulk of the interfaces will be externally facing RJ45 interfaces and will be used for both LAN (user facing) and WAN connections.

LAN connections will range from 8 – 24 ports

WAN connections will range from 2 – 4 ports

SFP/SFP+ Ports

SFP and SFP+ Ports will be used for multiple purposes:

- WAN connections (Both 10Gbps and 1Gbps)
- Downlinks to customers network
- East/West links between Open CPE platform devices when multiple are deployed to a site

SFP+ ports must be backward compatible to 1G SFP modules.

Power over Ethernet

The Ethernet ports should support Power over Ethernet to allow the attachment of Wifi AP's, Telephones or other devices. The PoE specification should be 802.3at at a minimum.

Internal Expansion Module

The Add-on Module slot allows for the addition of optional components such as LTE or Wifi, or other devices.

The form factor should be a standard M.2 connector/board design. M.2 provides support for PCI, SATA and USB interfaces to the host.

Redundant Power Supplies

The systems should have the option to support redundant power supplies, as well as the ability to detect and send a “dying gasp” or “last gasp” alarm when a power loss is detected.

For cost reasons, lower end platforms may not have a redundant power supply.

System Specifications

This section of the document outlines in more specific detail various models of the Open uCPE that AT&T is

putting forth to the OCP Community. The models range from “Micro” up to “Large” and vary based primarily on performance requirements.

Common Requirements

SFP Support	<p>The devices should support the use of 3rd party SFP's.</p> <p>The SFP+ ports should be backward compatible to SFP.</p>
LED Status Indicators	<p>The device should have LED lights incorporated that show at a minimum the following:</p> <ul style="list-style-type: none"> Power (for each power supply if multiple) System (Solid, Blinking) Alarm (Red, Amber) System Locator (Blue blinking) Link Status Link Speed Link Duplex <p>Note: Link status may be represented by same LED through use of a Mode button.</p> <ul style="list-style-type: none"> Phone Home Complete (Green, Blinking Green, Amber, Blinking Amber) <p>The System, Alarm and Phone Home indicators need to be controlled via AT&T software.</p>
Factory Reset	<p>The device should support a recessed Factory Reset button or sequence that will cause the system to revert to a startup configuration as if it came from the factory.</p>

Hardware Sizes

Micro

Hardware Specs	<p>CPU: 2 Core (Atom C3338) Memory: 4GB DDR4 ECC Storage: 32GB SSD LAN: 2 1GE RJ45 WAN: 2 1GE RJ45 & 2 SFP (Combo RJ45/SFP Ports) Host Mgmt: 1 GE RJ45 USB: 2x Front, 2x Rear, 1x Internal horizontal Serial Console: 1 LTE: Optional built-in module with externally accessible SIM and antenna Wifi: Optional built-in module and antenna M2 Expansion slot: Additional M2 connector supporting up to 30x110mm card for future module support and antenna connector BMC: Optional TPM: Yes, TPM 2.0 Power: Single power supply with Dying Gasp</p>
Hardware Form Factor	The Extra Small device should be a small form factor clamshell design with desktop mount as the default. The ability to mount in a rack with a bracket kit or on a wall (with wall mount kit) is optional but preferred.

Extra Small

Hardware Specs	<p>CPU: 4 Core (Atom C3558) Memory: 8GB DDR4 ECC Storage: 64GB SSD, 16GB SSD (Recovery Drive) Switch Chip: Marvell 88E6190 Internal Switch Ports: 2x 2.5GE External Switch Ports: 6x 1GE RJ45, 2x SFP PoE: 2 802.1at PoE+ ports Host Mgmt: 1 GE RJ45 (Shared with BMC) USB: 2x Front, 2x Rear, 1x Internal horizontal Serial Console: 1 LTE: Optional built-in module with externally accessible SIM and antenna Wifi: Optional built-in module and antenna M2 Expansion slot: Additional M2 connector supporting up to 30x110mm card for future module support and antenna connector BMC: Yes, OpenBMC TPM: Yes, TPM 2.0 Power: Single power supply with Dying Gasp</p>
Hardware Form Factor	The Extra Small+ device should be a small form factor clamshell design with desktop mount as the default. The ability to mount in a rack with a bracket kit or on a wall (with wall mount kit) is optional but preferred.

Small

Hardware Specs	<p>CPU: Xeon-D 4 Core Memory: 16GB DDR4 ECC Storage: 200GB SSD, 32GB SSD (Recovery Drive) Switch Chip: Broadcom Hurricane3 BCM56160 Internal Switch Ports to Host: 4x 1GE External Switch Ports: 10x 1GE RJ45, 4x SFP PoE: 4 802.1at PoE+ ports Host Mgmt: 1 GE RJ45 (Shared with BMC) USB: 2x Front, 2x Rear, 1x Internal horizontal Serial Console: 1 LTE: Optional built-in module with externally accessible SIM and antenna Wifi: Optional built-in module and antenna M2 Expansion slot: Additional M2 connector supporting up to 30x110mm card for future module support and antenna connector BMC: Yes, OpenBMC TPM: Yes, TPM 2.0 Power: Dual Redundant with Dying Gasp</p>
Hardware Form Factor	The device should be a 1RU rackmount design.

Medium

Hardware Specs	<p>CPU: Xeon-D 8 Core Memory: 32GB DDR4 ECC Storage: 400GB SSD, 64GB SSD (Recovery Drive) Switch Chip: Broadcom Hurricane3 BCM56160 Internal Switch Ports to Host: 2x 10GE, 2x 1GE External Switch Ports: 10x 1GE RJ45, 2 SFP+, 2 SFP PoE: 4 802.1at PoE+ ports Host Mgmt: 1 GE RJ45 (Shared with BMC) USB: 2x Front, 2x Rear, 1x Internal horizontal Serial Console: 1 LTE: Optional built-in module with externally accessible SIM and antenna Wifi: Optional built-in module and antenna M2 Expansion slot: Additional M2 connector supporting up to 30x110mm card for future module support and antenna connector BMC: Yes, OpenBMC TPM: Yes, TPM 2.0 Power: Dual Redundant with Dying Gasp</p>
Hardware Form Factor	The device should be a 1RU rackmount design.

Large

Hardware Specs	<p>CPU: Xeon-D 16 Core Memory: 64GB DDR4 ECC Storage: 1x 100GB SSD, 2x 3TB 7200RPM HDD Switch Chip: Broadcom Hurricane3 BCM56172 with 54998E 2.5G phy's Internal Switch Ports to Host: 2x 10GE, 2x 1GE External Switch Ports: 18x 2.5GE RJ45, 8 SFP+ PoE: 16 802.1at PoE+ ports Host Mgmt: 1 GE RJ45 (Shared with BMC) USB: 2x Front, 2x Rear, 1x Internal horizontal Serial Console: 1 LTE: Optional built-in module with externally accessible SIM and antenna Wifi: Optional built-in module and antenna M2 Expansion slot: Additional M2 connector supporting up to 30x110mm card for future module support and antenna connector BMC: Yes, OpenBMC TPM: Yes, TPM 2.0 Power: Dual Redundant with Dying Gasp</p>
Hardware Form Factor	The device should be a 1RU or 2RU rackmount design.

Overview of Sizes

	Micro	Extra Small	Small	Medium	Large
CPU	2C Atom	4C Atom	4C Xeon-D	8C Xeon-D	16C Xeon-D
Memory	4GB	8GB	16GB	32GB	64GB
Storage	32GB SSD	16GB SSD 64GB SSD	32GB SSD 200GB SSD	64GB SSD 400GB SSD	100GB SSD 2x 3TB HDD
Switch Chip	None	Marvell	Broadcom	Broadcom	Broadcom
Internal Switch Ports	None	1x 2.5GE	4x 1GE	2x 10GE 2x 1GE	2x 10GE 2x 1GE
External Switch Ports	None	6x 1GE RJ45 2x SFP	10x 1GE RJ45 4x SFP	10x 1GE RJ45, 2x SFP+ 2x SFP	18x 2.5GE RJ45 8x SFP+
PoE Ports	None	2x 802.1at	4x 802.1at	4x 802.1at	16x 802.1at
Host External NIC's	4x 1GE RJ45 2x SFP	None	None	None	None
Host Mgmt Port	1x 1GE RJ45	1x 1GE RJ45	1x 1GE RJ45	1x 1GE RJ45	1x 1GE RJ45
USB Ports	5x USB 3.0	5x USB 3.0	5x USB 3.0	5x USB 3.0	5x USB 3.0
Serial Console	Yes	Yes	Yes	Yes	Yes
LTE	Optional Module	Optional Module	Optional Module	Optional Module	Optional Module
Wifi	Optional Module	Optional Module	Optional Module	Optional Module	Optional Module
M2 Expansion	Yes	Yes	Yes	Yes	Yes
BMC	Optional	Yes	Yes	Yes	Yes
TPM	Yes	Yes	Yes	Yes	Yes
Power Supply	Single, w/Dying Gasp	Single, w/Dying Gasp	Dual, w/Dying Gasp	Dual, w/Dying Gasp	Dual, w/Dying Gasp

Software Support

The Open CPE Network Compute Platform supports a base software package composed of the following components:

BMC support

OpenBMC

ONIE

Initialization and firmware updates

Switch Abstraction Interface (SAI)

Hardware independent switch programming API

Open Network Linux

See <http://opennetlinux.org/> for latest supported version

Specification Requirements

Environmental placement requirements differ for the outdoors and indoors units.

Safety

- UL/ Canada
- CB (Issued by TUV/RH)
- China CCC

Electromagnetic Compatibility

- GR-1089-CORE
- FCC Title 47, Part 15, Subpart B Class A

ROHS

Restriction of Hazardous Substances (6/6)

Compliance with Environmental procedure 020499-00 primarily focused on Restriction of Hazardous Substances (ROHS Directive 2002/95/EC) and Waste and Electrical and Electronic Equipment (WEEE Directive 2002/96/EC)