AT&T Open Programmable-PON
1RU OLT Specification

Revision 1.0

Author: Sumithra Bhojan

Acknowledgement

I would like to acknowledge the contribution and partnership of Cambridge team led by Fazil Taslimi and Ali Taslimi in development of this spec. Also to my colleagues at AT&T and at other Service Providers who provided invaluable insights into how the Open FPGA could be used. Lastly, the OCP, without whose support and existence this type of work would not be possible.

-Sumithra

Revision History

Revision	Date	Author	Description
1.0	09/22/17	Sumithra Bhojan	Initial Draft

Contents

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 Programmable-PON 1RU OLT submitted to the Open Compute Project.

Overview

This document describes the technical specifications of the AT&T Open Programmable-PON 1RU OLT using FPGA. The system is a physical 1RU that is self- contained, and not part of a physical chassis system. However, like the GPON system, this physical box is attached to a fabric (e.g. leaf and spine) that interconnects many to create large scale-out virtual network elements[1]. The Open Programmable-PON 1RU OLT is a high-performance access design focused on NFV Infrastructure deployments which support 10Gbps/25Gbps PON access connectivity and provide up to 100Gbps uplinks to the ToR (Top of Rack) or Spine layer of the network. This design is intended to future proof the hardware design to include current XGSPON/NGPON2 and 25G PON readiness. One of the key benefit is that the FPGA’s can be dynamically

[1] This type of system is described in ETSI NFV architecture, where it comprises the Infrastructure that supports Network Functions Virtualization – often called NFVI. Additionally, the open software beyond that described in this specification is collected and distributed by the Linux Foundation as part of the CORD project: http://opencord.org

configured to support higher speeds as subscriber's transition to the next gen PON.

The switch supports 24x XFP PON ports that each operate at 10G/25G downstream (egress) and 10G/25G upstream (ingress) and six 100G QSFP ports. Other applications are also envisioned and described below.

The Open Programmable-PON 1RU OLT is a PHY-Less design with the XFP connections directly attaching to the XFI and SERDES interfaces of the xPON FPGA for OLT. The Open Programmable-PON 1RU OLT supports traditional features found in Top of Rack switches such as:

- Redundant field replaceable power supply and fan units
- A variety of supported power supply voltages
- Support for "Front to Back" or "Back to Front" air flow direction
- A 1RU design that supports standard 19" rack deployments as well as standard 21" Open Rack and also 23" telco rack deployments.

Applications of this design include some specific variations and considerations in order to address a broader market and set of collaborators for this project. Specifically, the set of variations considered include:

1. The line card application for a virtualized scale-out OLT.
2. A small, simplified, self-contained traditional OLT device that can be used in non-NFVI environments.
3. The ability to manage and control the device using a separate out-of-band port.
4. The ability to synchronize and distribute timing from upstream toward ONUs.

The four variations are not mutually exclusive, and are now described in more detail.

Virtual OLT

The first application embodies the virtual OLT. In this application the device is envisioned to be part of a NFVI deployment, where compute, storage, and other cloud infrastructure are part of the environment. In this environment, the Open Programmable-PON 1RU OLT connects to leaf or spine devices that aggregate traffic and also provide transport for management and control.

The low-level management processes that might be typically performed by an embedded processor on a Programmable-PON 1RU OLT are performed outside the box, in a separate commodity server. This allows managing many OLT devices from a small number of commodity servers, and conserves compute and storage in the same way that performing aggregation in the ToR switch conserves aggregation typically performed on line cards to interface several 10/25G-PON PHY chips to a backplane or fabric. The arrangement also allows for an overall system design where there is no single point of failure that affects the entire Open Programmable-PON 1RU OLT and allows application of cloud application architecture to software that was previously hosted on embedded processors.

This system is described as a disaggregated or virtual OLT because the various software and hardware components that were once integrated into a single physical device have been separated and supported in a distributed way across NFV infrastructure. By doing this it becomes more likely to re-use components among disparate networking functions and to independently scale resources and investments according the specific application of the technology.

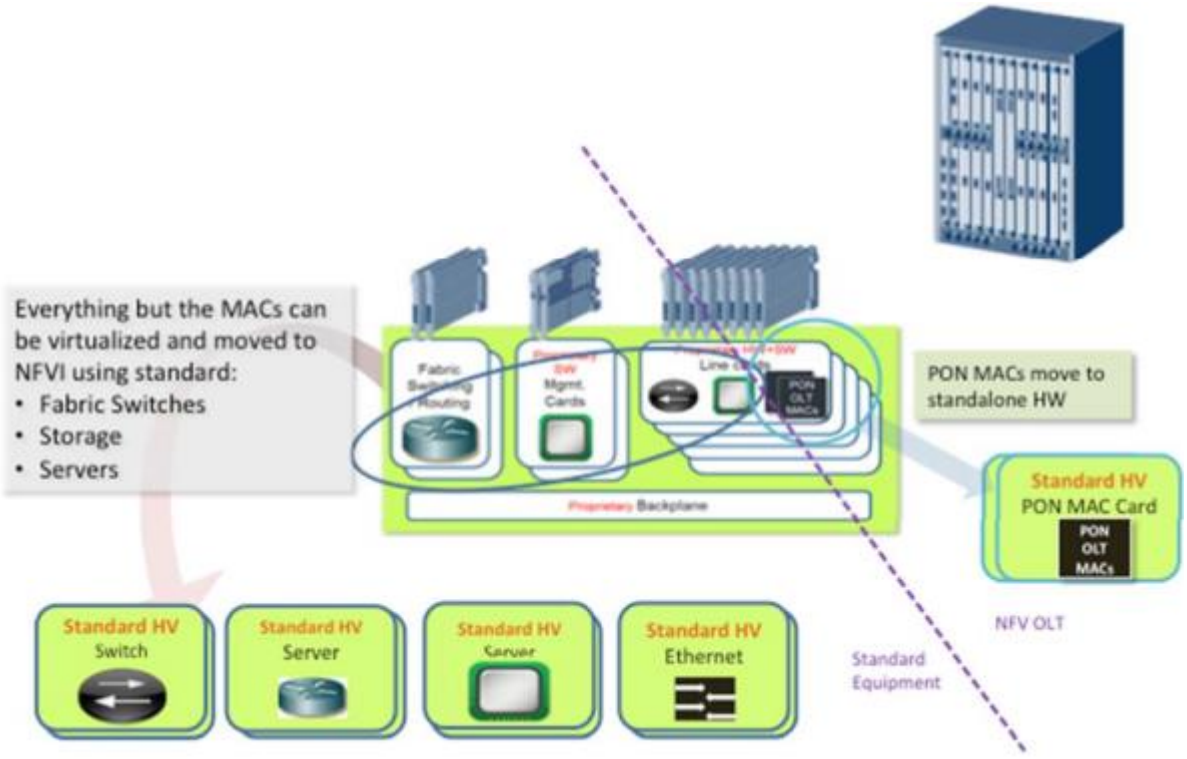

Figure 1 - Virtual OLT

Figure 1 shows the disaggregation of a typical OLT and the mapping of its functions to NFV infrastructure. The backplane or fabric of the OLT is mapped to the fabric of NFVI (ToR and EoR / leaf & spine switches). The interfaces to the fabric (what were card edge connectors) become simple Ethernet interfaces. The control and management functions are embodied in SDN control and Orchestration with a software stack that will be described shortly. Processing and configuration are mapped to compute and storage in NFVI. The figure shows that almost all the components in an OLT, those to the left of the dashed line, can be mapped to standard high-volume (HV) components found in NFVI. The exception is in the FPGA-PON PHY and MAC. Those are not typically found in NFVI and this specification describes a standard HV device that supports them.

In the Virtual OLT, the application of the AT&T Open Programmable-PON 1RU OLT is to facilitate attaching OLT silicon to the fabric of NFVI – as shown in Figure 2.

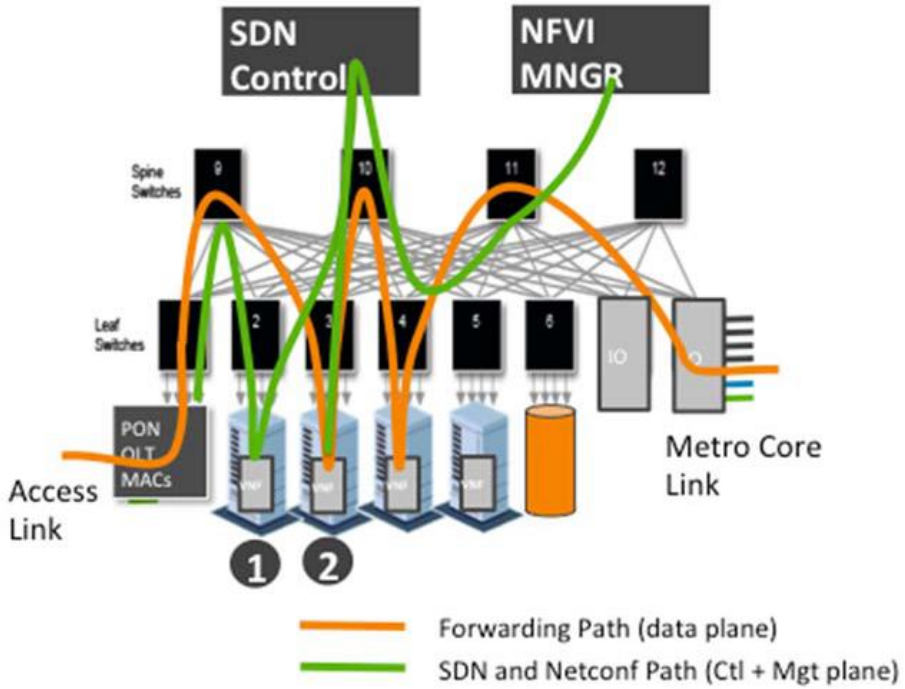

Figure 2 - Typical application of XGS-PON 1RU OLT

Figure 2 shows the Open Programmable-PON 1RU OLT attached to typical ToR (Leaf) switches in the lower left corner. Shown prominently at the top are processes that are run inside the compute VNFs at point 1 and 2. These processes provide the management and control plane functions that manage and control the overall system including the FPGA-PON silicon within the Open Programmable-PON 1RU OLT.

The software used to support this system is partly shown in Figure 3. At the lowest level of the figure we find firmware and hardware drivers that are part of the software loaded and run on the FPGA chip. That software is loaded on the chip and then subsequently configured and managed through a matching low-level API along with an OMCI stack which is run on a commodity data center (DC) server. These elements are colored blue because they are proprietary and specific to the silicon, and would need to be replaced to support other chips and PHY technologies.

The next layer up is shown in purple and represents open source software. The lowest layer of this software consumes OLT API calls and OMCI signaling and creates a homogenized abstraction of an OpenFlow controlled OLT – largely patterned after an Ethernet switch. This abstraction is then plugged into the southbound side of an OpenFlow agent and configuration management block, and that software comprises all the code to generically manage and control the Open Programmable-PON 1RU OLT.

The next layer in the software stack is an OpenFlow Controller and configuration. These might be combined or separate software. Finally, in green, we see the applications that embody the control plane and management applications for a number of access technologies and instances of each.

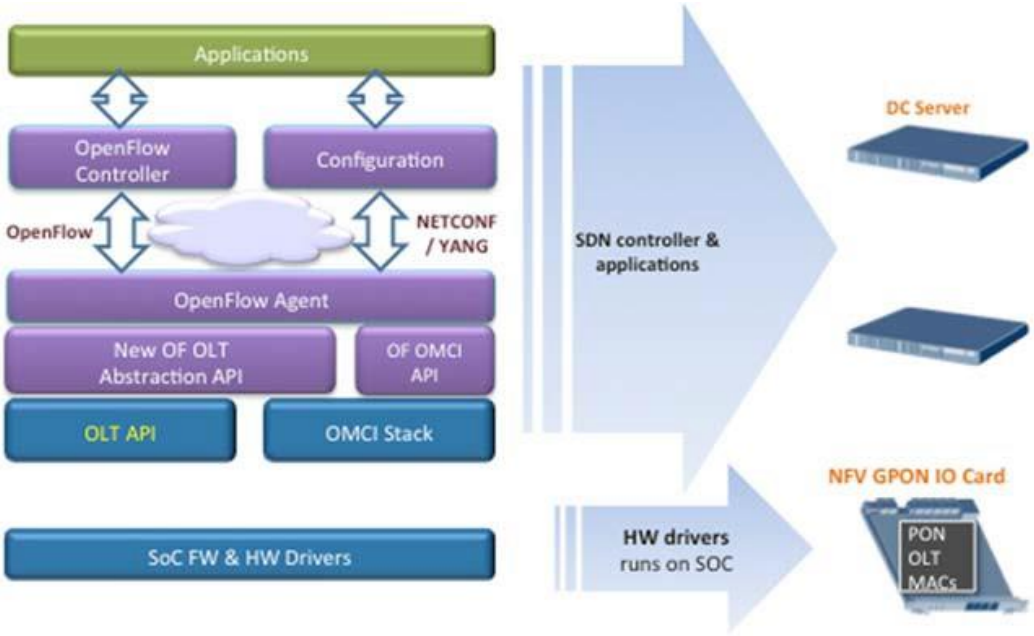

Figure 3 - Software stack for Open Programmable-PON OLT

To support this application, the Open Programmable-PON 1RU OLT minimizes the number of components populated on the system board and leverages external, scalable, available and re-usable components instead.

This is the most basic and economic instantiation of this design. Management is performed through external processes and is communicated through a virtual LAN that isolates that traffic from customer traffic. Specifically, there is a need to support OOB (out of band management) LAN. The design does still support a variety of power supply options, airflow options, and a baseboard management controller (BMC) to manage these resources.

Mini OLT

The next application of this technology moves to the opposite extreme. In this case the Open Programmable-PON 1RU OLT is extended with Processing, Storage, and an internal PCIe communications to facilitate the device becoming a stand-alone OLT system. From a hardware perspective, the board is populated in addition to the components of the virtual OLT with a CPU and RAM, and Flash storage. This application is useful in one-off applications, in lab and development environments, and in places where NFVI isn't available or desired. The software stack for this device can be identical to that used in the virtual OLT: where either part or even all the software shown in **Error! Reference source not found.** is hosted on the onboard CPU. However, it is also possible to optimize. For example, control and management applications could be written directly to the merchant silicon APIs. Obviously, hybrid approaches are also possible by making other software placement choices, and other software not shown above could also be run on the CPU. Note however, that this design has several single points of failure, and is not as robust from the perspective of availability as the virtual OLT.

Timing

The next application covers a case for timing distribution. This application is largely independent of the previous ones. That is, regardless of whether there is an onboard software, processor and storage, there are situations when distributing timing is desired. The typical case for including timing distribution is where XGS-

PON/NGPON2/25G PON are used to provide backhaul for cell sites, but this is not a universal requirement for using 10G or higher PON for cell site backhaul. In this application an optional Synchronization Management Unit is added to the system to facilitate SyncE and IEEE 1588 timing paths.

This overall design supports all the aforementioned use cases, and allows the omission or depopulation of various components as manufacturing options to source the device for the various use cases.

XGS-PON Network

A 10 Gigabit Services PON (XGS-PON) network supports symmetrical 10G/10G communications to enable a variety of applications. Some key features include:

- Data rates of : ~10Gbps (9953Mbps) DS/~10Gbps (9953Mbps) US
- At least 128 ONTs (downstream PON endpoints) per PON Link
- ToD synchronization
- ONT power saving operation
- Optional downstream AES encryption for each port
- IPv6
- Dual-stack IPv4 and IPv6
- Symmetric bandwidth allocation

Figure 4 shows a typical XGS-PON Network.

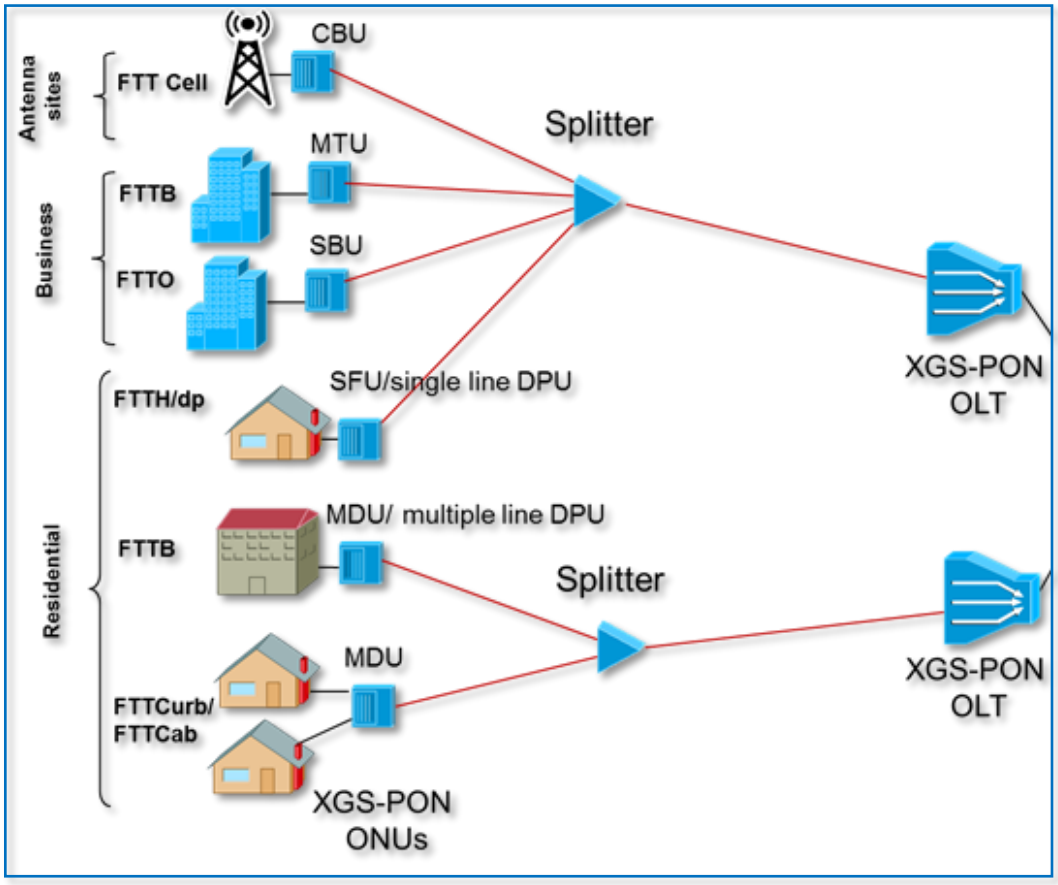

Figure 4 - XGS-PON Network

System Overview

An overview of the system in terms of functional blocks is seen in Figure 5.

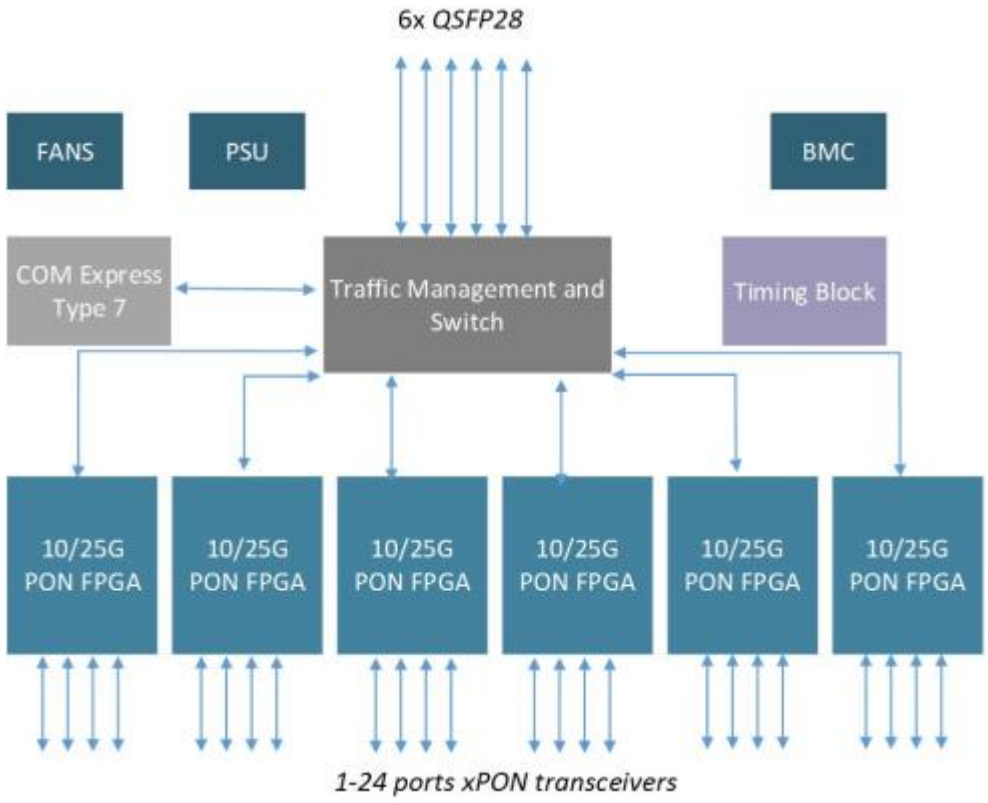

Figure 5 - Main System Block diagram

The Open Programmable-PON 1RU OLT system is a physical 1RU box and consists out of the following functional hardware modules:

10/25G PON FPGA

A high-performance, cost effective OLT PON MAC FPGA with support for XGS-PON, NGPON2 and future 25G PON technology. It provides up to four 10/25G interfaces. Six FPGA's are used to support 24 ports of 10/25G PON in this specification.

A high-performance family of FPGA is used for this application. This FPGA family of products integrates a feature-rich 64-bit quad-core or dual-core ARM based processing system (PS) and programmable logic (PL) architecture in a single device. The integrated quad-core 64bit ARM is mainly used to offload the main processor. Functions such as DBA, PLOAM, OMCI, PON management, and etc. Two high performance integrated PS cores are for OMCI message handling and the other interactive operations.

Host CPU

The Host Module will use standard COM-Express type 7 form factor and interface. This supports choosing from a variety of different CPU options with different performance and price points including Intel Rangeley, D-1500, and ARM. COM-Express allows for different form factors in a compatible mounting strategy which this platform will support. A typical Host CPU for stand- alone OLT might include D-1500 2 to 4 Core with 16G RAM and 64G M.2 SSD.

BMC

- Serves for general box management
  - For example: I2C control, LEDs, Interrupts + I/O expender, Fan and PS control, Power sequencing
- Open BMC

NPU

NPU performs the functions of PON Traffic management and Switch. This is optimized to provide cost-effective seamless 1GBASE-T to 10GBASE-T and 10GbE to 25GbE access with support for six 40GbE/100GbE uplink connectivity.

Physical Overview

Maximum Dimensions

	Inches	Millimeters
Length -	24	609.6
Width – 19” rack	17.5 with ears for 19”	444.5
Height – 1 RU	1.75	44.45
Note: Width does not include mounting ears – which must have holes or closed slots. Depth does not include PSU handles.		

Top View

The top view of the Open Programmable-PON 1RU OLT shows the PCBs and associated components in the system (illustrative only).

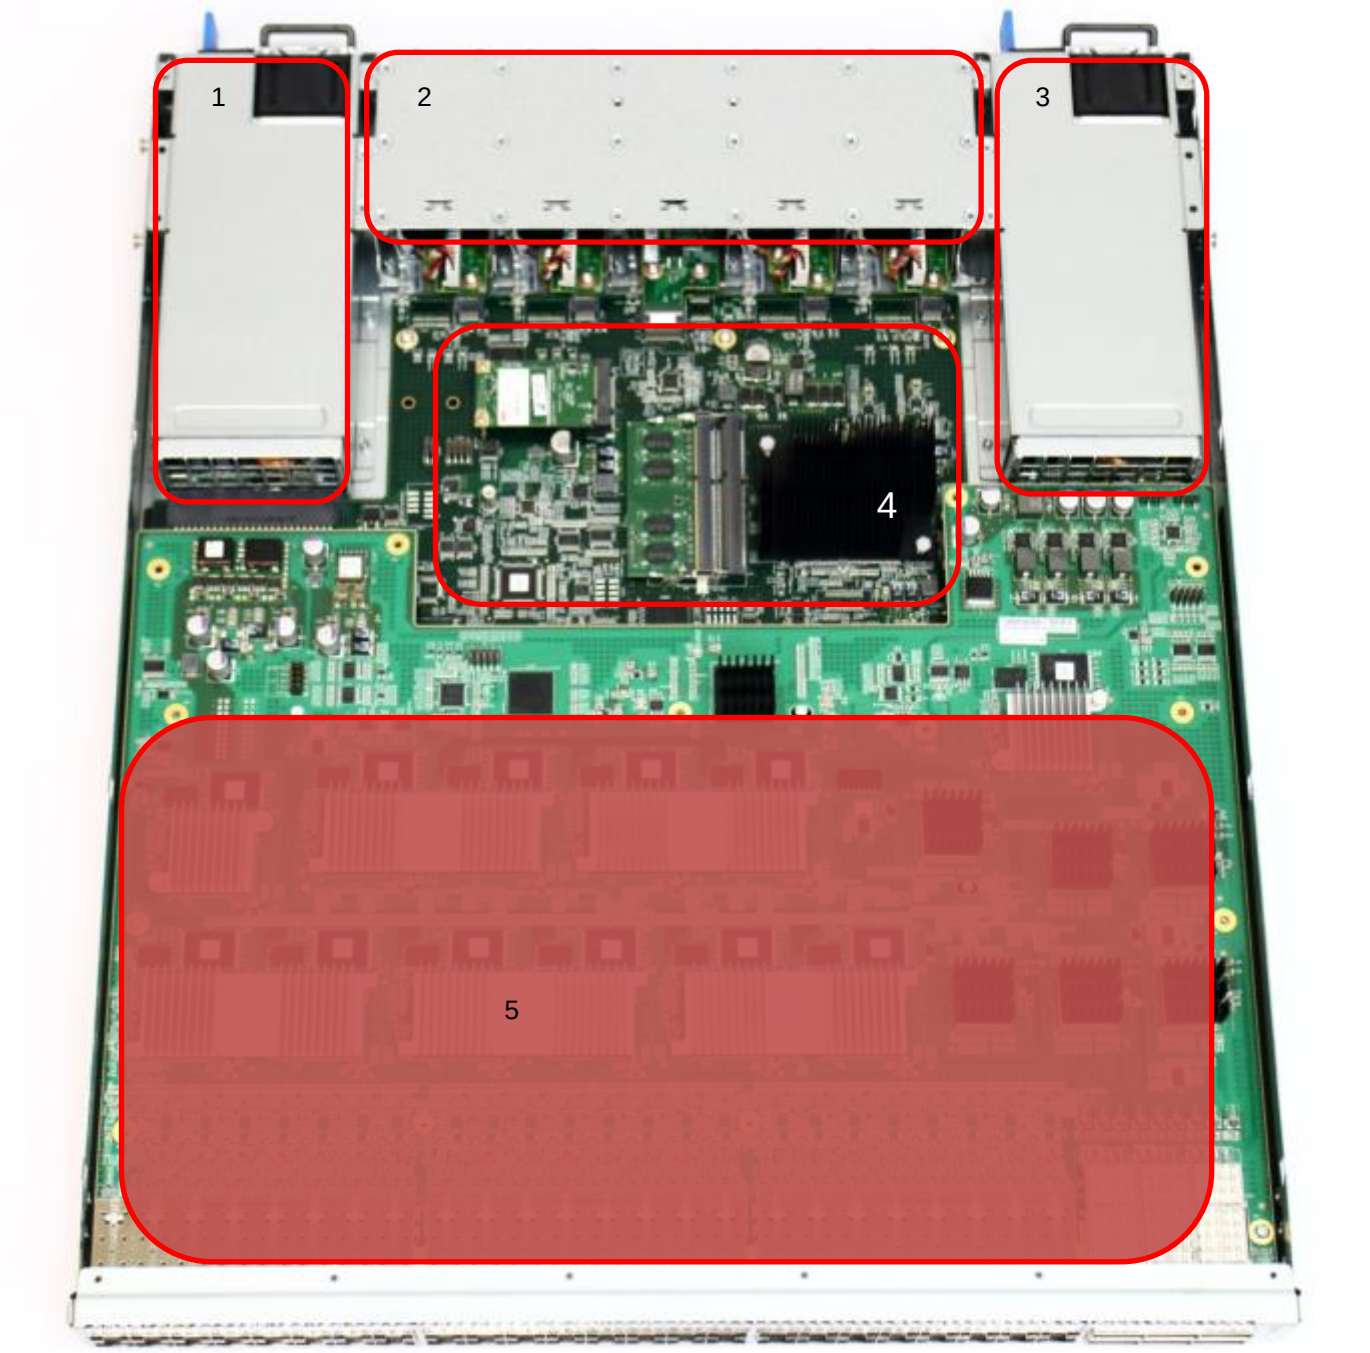

Figure 3 - Top View (Illustrative)

1. Hot swappable power supply #1
2. Hot swappable fan modules
3. Hot swappable power supply #2
4. CPU system
5. 6 PON FPGA devices and 1 aggregating switch

Front View

Figure 7 - Front Panel (Illustration only)

- PON Links: 24x XFP ports
- Uplink Ports: 4x QSFP

- LEDs
    - XFP Module status
    - QSFP port status
    - System and PSU LED indicators
    - System Finder LED

Rear View

- Five (4+1) redundant hot swappable fan modules (Including Color coding to indicate airflow direction)
- Two redundant hot swappable power supply modules
    - LED per power supply to indicate status
    - Color coding to indicate airflow direction

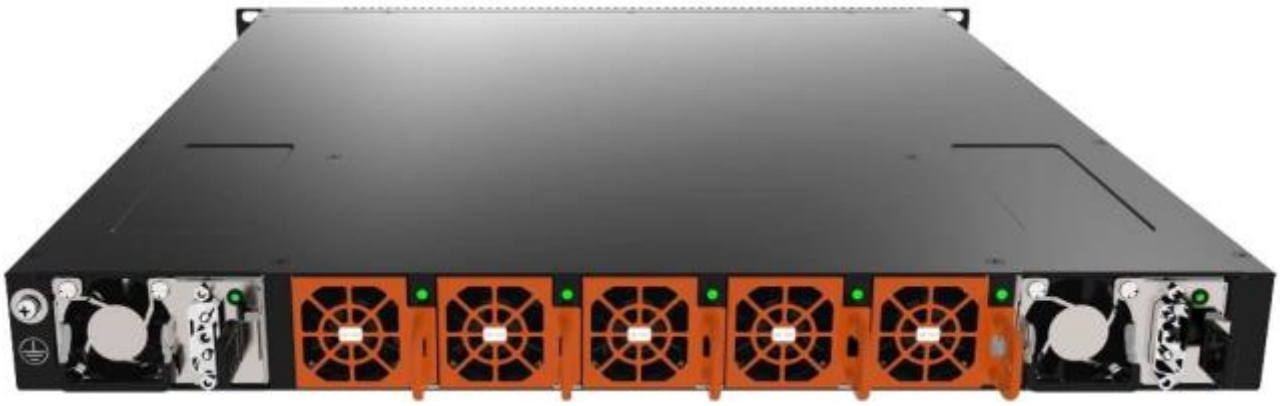

Figure 8 – Rear view (Illustration only)

Panel LED Definitions

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
Diag	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 Off – Function not active
XFP LEDS	LED built into XFP cage to indicate port status	On Green/Flashing – Port up with active ONTs Flashing indicates activity On Amber – Port up with no active ONTs Off – No Link/Port down
QSFP Break out LEDs	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?) Flashing indicates activity Off – No Link
OOB LED	LED to indicate link status of 10/100/1000 management port	On Green/Flashing - port has link Off – No link

Field Replaceable Units

This section describes field replaceable modules used in the design. These include optical XFP and QSFP transceivers, fans and power supplies. All modules must be hot swappable without the use of tools.

Power Supply Modules

The Open Programmable-PON 1RU OLT supports two hot swappable Power Supply Modules (PSMs) and needs only one to operate. Please use the below as a general guidelines for the PSUs selection:

- The PSM form factor should re-use an existing OCP form factor. (optional)
- The PSM must meet the power requirements of the design: e.g. 640W or more.
- The system must accept different PSM types (e.g. AC or DC) which have the same form factor.
- All PSMs must be available with F2B and B2F airflow.
- AC PSMs must support AC input between 100 & 240 VAC.
- AC PSMs must have a mechanism that prevents accidental dislodging of the cord.
- DC PSMs must support DC input between -57 & -40VDC.
- DC PSMs have additional physical connection requirements listed in AT&T TP 76450, section 2.4.

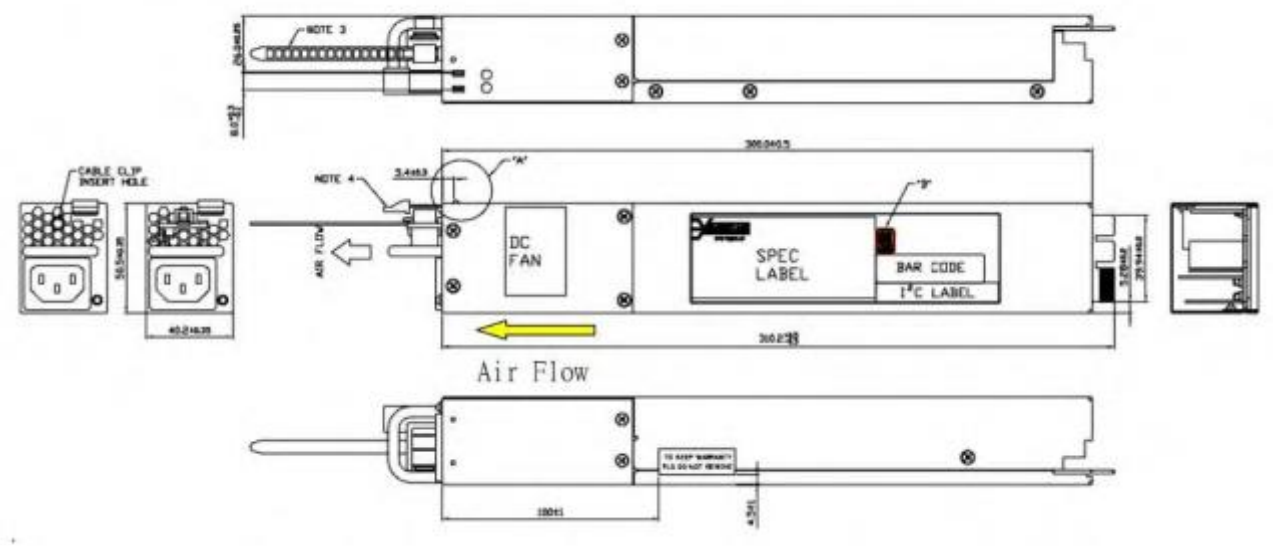

Figure 4 – Power Supply (Illustrative)

Fan Modules

The Open Programmable-PON 1RU OLT supports multiple individual fan modules in an N+1 scheme where N modules are sufficient to cool the device. Fan modules may use different designs for number and size of fan, as long as the overall design cools appropriately with a failed device. In 1RU designs, a typical fan module might have two 40mmx40mmx54mm fans. An example of a suitable fan module is show from the front and back in Figure 5. There are two fan module types differentiated by airflow, Front-to-back and back-to-front (F2B and B2F).

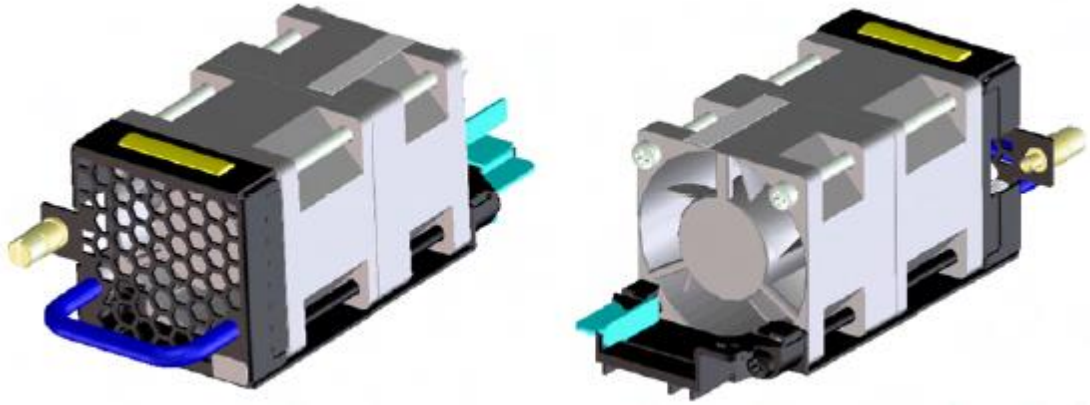

Figure 5 – Fan (Illustrative B2F)

The design should re-use fan modules typically found in existing OCP specs.

PCB Board Set

The Open Programmable-PON 1RU OLT is composed of several modules or PCBs. The design must have a separate board for the CPU to allow for different build options and for the CPU board to be omitted from a build completely.

The design may have additional modularity in PCB layout. For example, it's acceptable to have a separate fan board module as a re-useable component across different system designs. Similarly so for other build

options, like timing modules, and standard components like a BMC module.

Preferred designs for this spec. will re-use existing modules from other OCP Accepted designs.

Software Support

The Open Programmable-PON 1RU OLT supports a base software package composed of the following components:

BMC support

AMI BMC or OpenBMC

U-Boot (Optional CPU Module)

TBD

ONIE (Optional CPU Module)

ONIE version 2014.08 or greater will be supported

Open Network Linux (CPU Module)

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

Overall System Software

See http://opencord.org for information and documentation

See https://github.com/opencord/cord for software source

Specification Requirements

AT&T has established specifications for servers and NFVI communications equipment that are intended to be used in Central Offices. These specs revisit classical NEBs requirements – particularly in the face of new resiliency and availability architectures. Servers that are not a single point of failure for their services (e.g. follow a typical cloud resiliency model) follow ATT-TP-76207. Telco devices that do have or comprise single points of failure for their services including this specification follow ATT-TP-76208. These latter requirements are partially repeated here.

Specifications that require testing must be confirmed by an accredited agency recognized by the National Cooperation for Laboratory Accreditation or ISO/IEC Guide 25 or ISO/IEC 17025.

Note: These specifications are limitations placed on any design. The actual performance of the AT&T Open GPON OLT meets or exceeds these specifications.

Power Consumption

The total estimated system power consumption must be specified in watts. This is based upon worst case power assumptions for traffic, optics used, and environmental conditions. Typical power consumption should also be provided, as well as heat dissipation.

- ATIS TEER (ATIS-0600015.2009) should be measured and provided (Preferred)
- SPECpower_ssj2008 can be substituted for ATIS TEER (Acceptable)
- US EPA Energy Star Certification is favored.
- Power terminations must be clearly labeled and fully protected with a non-metallic, non-flammable cover.
- ATT-TP-76208 lists additional power requirements for under and over voltage, grounding, and current characteristics.

Environmental

- Light weight is favored
- 15 to 40°C operating range – de-rated 1°C for every 1000 ft (300 m) above 6000ft (1830m).
- Humidity 5% to 85% non-condensing (operational and storage)
- Vibration – IEC 68-2-36, IEC 68-2-6
- Shock – IEC 68-2-29
- Acoustic Noise Level – Under 78dB in 26 degree C
- Altitude: -200ft (-60 meters) to 6000ft (1830 meters).

Safety

Fire Spread. Field conditions for telco deployment may require deployment in existing Carrier Communications Spaces that utilize Fire Code Exemptions and do not have automatic fire suppression. NFVI equipment, like the AT&T Open GPON OLT deployed in these locations must meet enhanced fire spread requirements:

Generally, the equipment must meet ATIS-0600319.2014 *Equipment Assemblies – Fire Propagation Risk Assessment Criteria* (see note below).

Note: Equipment may conform to this requirement by way of inherent design features that include all four items below:

1. Height of 2 RU or less
2. Horizontally mounted main printed circuit board

3. Metallic 5 sided enclosure with a metallic or non-metallic front cover or faceplate
4. Non-metallic materials shall comply with ATIS-0600307 4.1

For equipment that does not meet the fire spread requirements of ATIS-0600319.2014 by way of inherent design features noted above, the manufacturer must attest that the equipment has successfully passed the burn test as referenced in the ATIS document.

- 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)