FiBr

TELECOM INFRA PROJECT

Open Fixed Access Networks

Technical Requirements Document

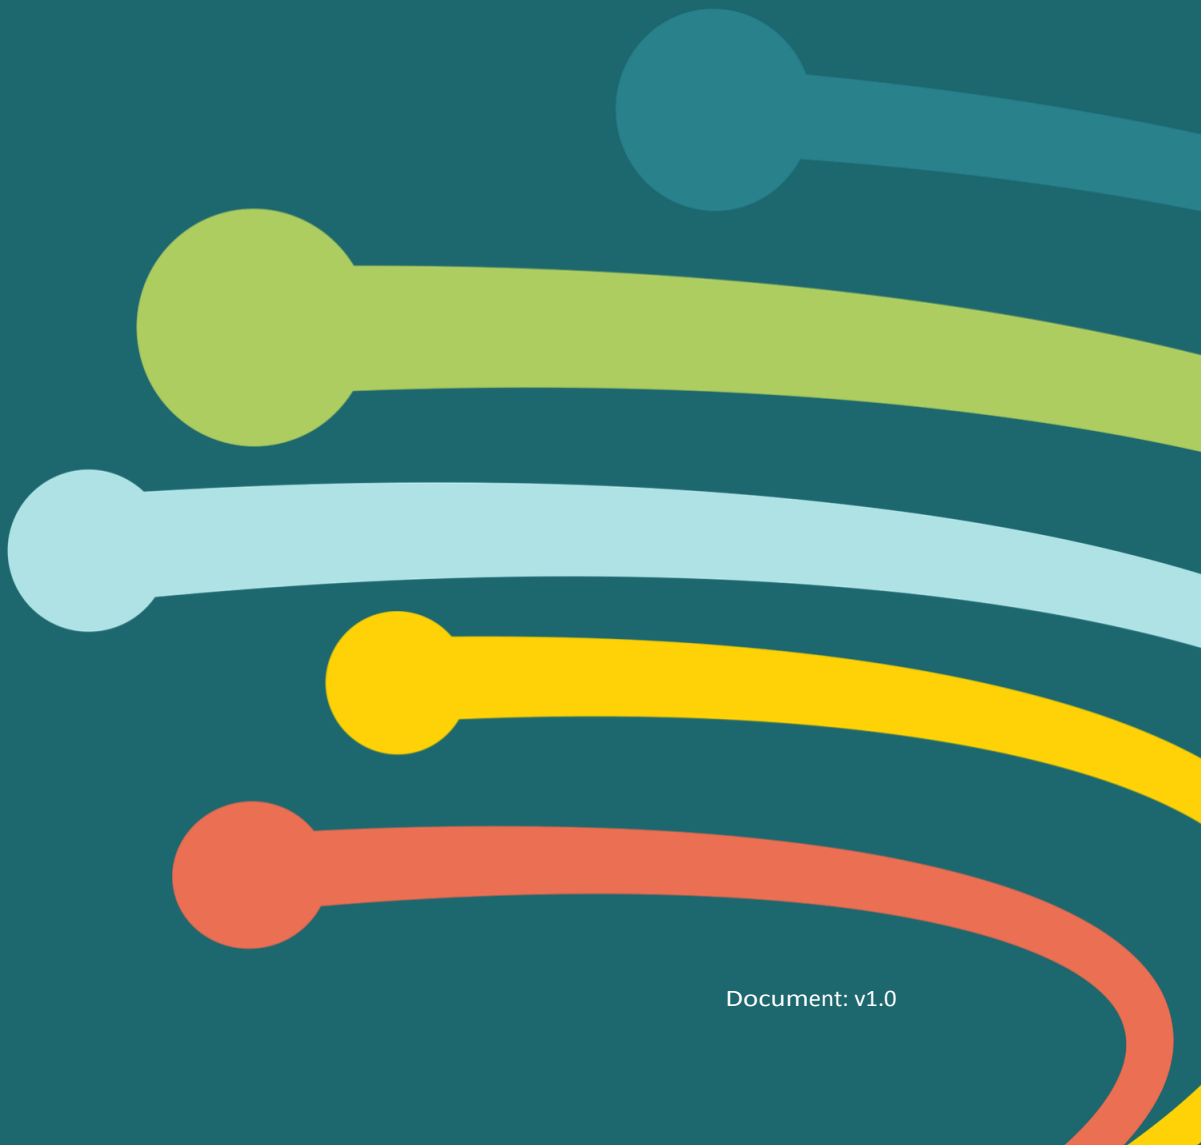

Date June 2023

Document: v1.0

Authors

Bruno Cornaglia

Fixed Access Senior Manager, Vodafone.

bruno.cornaglia@vodafone.com

Daniel Cortes Olmeda

Fixed Access Network Manager, Telefonica GCTIO.

daniel.cortes@telefonica.com

Ian Cooper

Senior Researcher, BT.

ian.r2.cooper@bt.com

Jose Torrijos Gijon

Technology Expert, Telefonica GCTIO.

jose.torrijosgijon@telefonica.com

Julio Montalvo Garcia

Technology Expert, Telefonica GCTIO.

julio.montalvogarcia@telefonica.com

Lorenzo Magnone

NGN Architect, Telecom Italia.

lorenzo.magnone@telecomitalia.it

Mauro Tilocca

Project Manager, Telecom Italia.

mauro.tilocca@telecomitalia.it

Paolo Pellegrino

Fixed Access Innovation Project Manager, Telecom Italia.

paolo.pellegrino@telecomitalia.it

Stefano Barbieri

Center Of Excellence Engineer, Vodafone.

Stefano.barbieri@vodafone.com

TIP Document License

By using and/or copying this document, or the TIP document from which this statement is linked, you (the licensee) agree that you have read, understood, and will comply with the following terms and conditions:

Permission to copy, display and distribute the contents of this document, or the TIP document from which this statement is linked, in any medium for any purpose and without fee or royalty is hereby granted under the copyrights of TIP and its Contributors, provided that you include the following on ALL copies of the document, or portions thereof, that you use:

1. A link or URL to the original TIP document.

2. The pre-existing copyright notice of the original author, or if it doesn't exist, a notice (hypertext is preferred, but a textual representation is permitted) of the form: "Copyright © 2022, TIP and its Contributors. All rights Reserved"

3. When space permits, inclusion of the full text of this License should be provided. We request that authorship attribution be provided in any software, documents, or other items or products that you create pursuant to the implementation of the contents of this document, or any portion thereof.

No right to create modifications or derivatives of TIP documents is granted pursuant to this License. except as follows: To facilitate implementation of software or specifications that may be the subject of this document, anyone may prepare and distribute derivative works and portions of this document in such implementations, in supporting materials accompanying the implementations, PROVIDED that all such materials include the copyright notice above and this License. HOWEVER, the publication of derivative works of this document for any other purpose is expressly prohibited.

For the avoidance of doubt, Software and Specifications, as those terms are defined in TIP's Organizational Documents (which may be accessed at https://telecominfraproject.com/organizational-documents/), and components thereof incorporated into the Document are licensed in accordance with the applicable Organizational Document(s).

Disclaimers

THIS DOCUMENT IS PROVIDED "AS IS," AND TIP MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE DOCUMENT ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

TIP WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE DOCUMENT OR THE PERFORMANCE OR IMPLEMENTATION OF THE CONTENTS THEREOF.

The name or trademarks of TIP may NOT be used in advertising or publicity pertaining to this document or its contents without specific, written prior permission. Title to copyright in this document will at all times remain with TIP and its Contributors. This TIP Document License is based, with permission from the W3C, on the W3C Document License which may be found at https://www.w3.org/Consortium/Legal/2015/doc-license.html.

Exhibit A

Draft Document Notice

TIP CONFIDENTIAL

Change Tracking

Date	Revision	Author(s)	Comment
09/12/2022	v1.0	OFAN sub-group	Initial version for release

Table of Contents

List of Figures

1

Introduction

1 Introduction

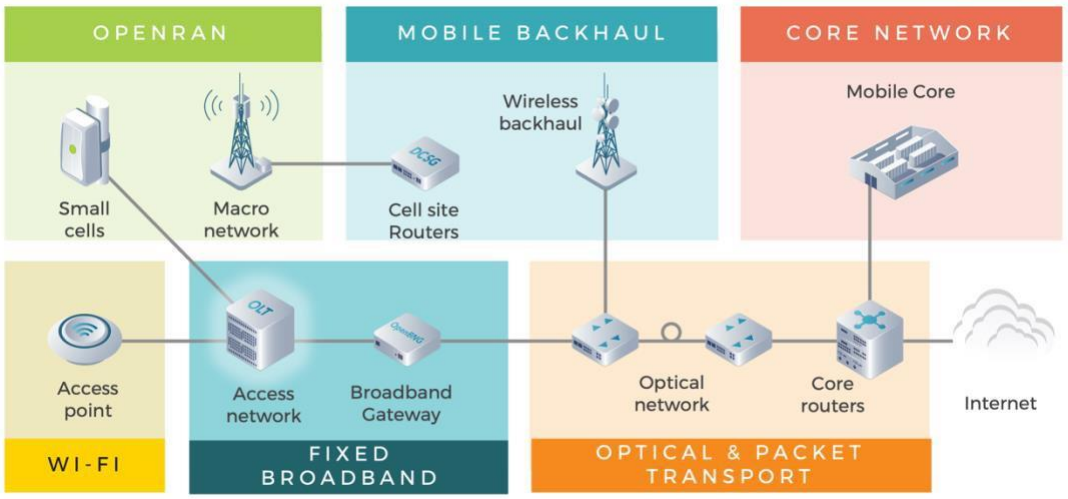

Figure 1 - TIP Fixed Broadband Project Group scope

The Telecom Infra Project (TIP) Fixed Broadband (FiBr) Project Group is working to develop a new generation of open and disaggregated technologies that help operators increase the availability of fast and reliable broadband services across the world. The Open Fixed Access Networks (OFAN) sub-group is part of the FiBr PG, tasked with the mission to build, test and deploy products that meet the needs of operators deploying access networks based on Passive Optical Network (PON) technologies.

This document describes the technical requirements for an open and disaggregated OLT device that can be deployed into a local exchange or central office environment. The document describes the required hardware variants and proposes software packages that can be combined with this hardware to enable provision of fixed access services.

1.1 Why Open Fixed Access Networks?

The ultimate objective for this work is to drive the deployment of interoperable access technologies. In so doing, the subgroup participants are looking to address the following topics:

- Multi-vendor interoperability, including north-south integration between OLT and management systems, and east-west integration between OLT and ONT
- Disaggregation, providing the ability to decouple the selection of HW from SW components to reduce the problem space for suppliers and increase supply chain diversity
- Programmability, enabling functional upgrades via modular SW packages; initially this is intended to simplify SW release/lifecycle management, but this may extend to cover integration of modules from multiple third-party suppliers over time
- Enable operational efficiencies, embracing new operational paradigms with more efficient devices that can be deployed into modern cloud-based technical architectures
- Unlock new service opportunities through the adoption of more distributed and converged network architectures

1.2 Scope of the document

The aim of this document is:

- To describe the OLT platform architecture and requirements that will need to be met by the system in terms of HW and SW features
- To describe the interfaces and associated requirements for other platforms and network elements that will integrate and interoperate with the OLT

Once this document is complete, the sub-group will then publish a detailed low-level technical requirements document. This document will describe requirements down to a feature level, and it will form the basis of onward discussions and collaborations with interested suppliers and other partners.

1.3 Document Structure

This document is structured as follows:

- **Platform architecture**, describes the overall system architecture and relevant overarching principles
- **Physical requirements**, characteristics for hardware (HW) products including 'speeds and feeds' and power/environmental optimizations

2

Platform Architecture

2 Platform Architecture

The OLT consists of a number of functions, which are provided by a combination of whitebox HW coupled with Network Operating System (NOS) SW. The diagram below depicts its high-level HW and SW functions, which will be described in further detail in later sections of this document:

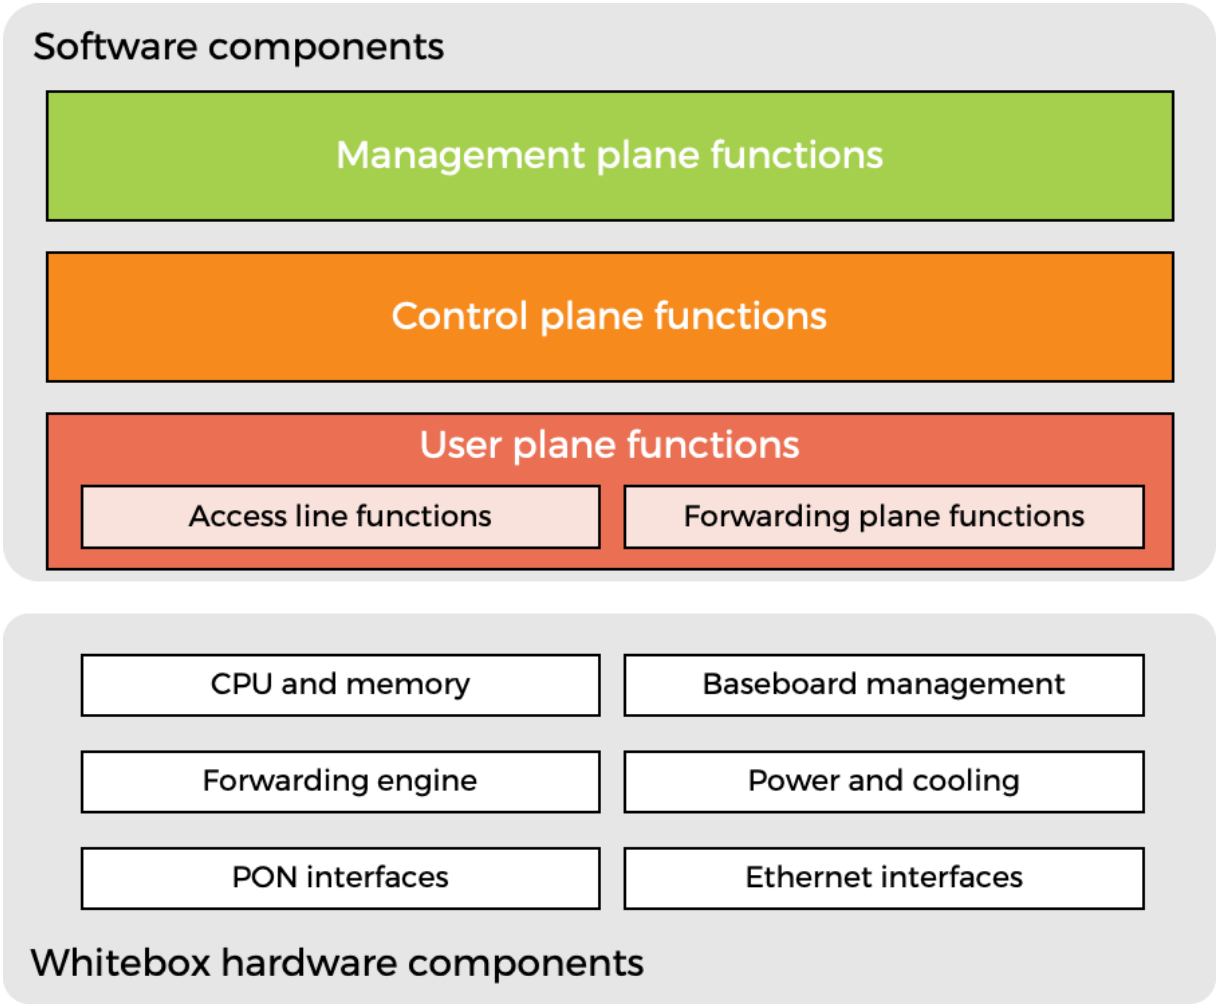

Figure 2 - OFAN platform architecture overview

For the Exchange-Optimized OLT use case, the OLT has a fixed ('pizza box') form factor. It is typically deployed into a Central Office environment, but may also be deployed into remote locations such as active street cabinets. These devices are optimized primarily for the delivery of FTTH and FTTP services, but may also be used to support small

numbers of additional services.

For further background regarding this architecture and the use case described in this document, please refer to the [OFAN-UCD] document.

2.1 Architecture principles

2.1.1 Disaggregation

One of the most fundamental tenets of this architecture is that it breaks down a traditional OLT 'monolith' into a number of discrete components. The intent here is twofold:

- It allows *suppliers* to focus on core/specialist competencies (e.g. hardware design, software engineering). This reduces the overall scope/problem space they need to address, and allows them to bring products and features to market more quickly
- It gives greater flexibility for *operators* over the technologies they use to build their networks. This allows them to select each component based on the metrics that are most relevant to them (e.g. cost, features, support, performance, energy consumption)

In the context of OFAN, disaggregation is primarily realized by the decoupling of hardware from software, with the ultimate goal of enabling 'mix and match' between different HW and NOS solutions. However, OFAN is also promoting the development of products and solutions that further disaggregate the SW stack. In turn, this offers up a more granular approach to deployment of OLT SW, with some functions provided locally on the OLT HW, and others provided in a local or remote cloud instance. To begin with, we expect that these disaggregated SW components would be provided as part of a single SW solution, but we welcome the possibility to apply the 'mix and match' principle to different SW elements as well.

While this document primarily focuses on these two interpretations/implementations of disaggregation, the OFAN group is open to additional forms of disaggregation, and would welcome feedback and contributions from the ecosystem on this topic.

2.1.2 Openness

The ‘open’ in Open Fixed Access Networks does not imply 'open-source', but rather that there must be no 'hidden' or 'restricted' APIs that preferentially benefit the hardware vendor; or the NOS vendor; or the OS selected. In particular, operators should be able to swap out or replace any component without affecting the other(s).

To this end, the hardware must not impose any restriction that limits the software that can run on it. In other words, it must be possible to install any operating system, even if its implementation comes from a third party. To ensure compatibility, it is highly recommended that Network Operating Systems for this platform are provided in the form of binary installers compatible with the Open Network Install Environment (ONIE) specification, and that OLT hardware is also equipped with ONIE.

Note that if the platform provides the capability to verify the signature (via a particular certificate or a cryptographic key) of the software, it must be possible to disable such verification at any time, through software or firmware configuration, without the need for any specific or additional license. Moreover the hardware must be accompanied with comprehensive documentation that provides maintenance, configuration, diagnostic and repair information.

2.1.2.1 Specific considerations for licensing of open solutions

OFAN is a technical activity, and the group does not intend to cover commercial aspects as part of this collaboration. However, the OFAN group feels that products developed according to these requirements should implement simple licensing structures that promote a high degree of transparency and predictability for the ecosystem, and where relevant these licenses should be focused on functionality rather than capacity, and should be implemented without requiring specific interaction with management and control infrastructure.

2.2 SW functions and reference interfaces

The OFAN solution is composed of a number of different SW functions, as depicted below:

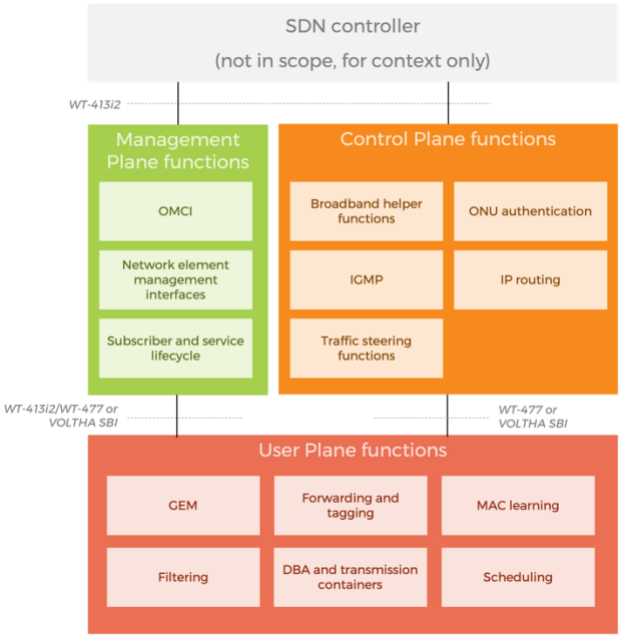

Figure 3 - SW functions and reference interfaces

The interfaces between SDN Controller and the Management Plane and Control Plane functions are NETCONF/YANG based and the data models are referenced in [WT-413i2].

The interface between the Management Plane and the User Plane functions can be either standard based according to [WT-413i2] / [WT-477] or compliant with established industry practices (e.g. [VOLTHA] SBI).

The interface between the Control Plane and the User Plane functions can be either standard based according to [WT-477] or compliant with established industry practices (e.g. [VOLTHA] SBI).

2.2.2 Cloudified architectures and SW modularity

In traditional OLTs all of these functions are provided locally, and each OLT is deployed and operated as a completely standalone network element. However, since a given operator may deploy many thousands of OLTs, operators increasingly see value in consolidating some of these functions in a cloud environment. This approach to disaggregation is known as CUPS (Control / User Plane Separation), and results in an Access Network that can be built and operated in a more flexible, nimble and sustainable way. An example implementation can be seen below.

Some OLT functions are located in the cloud (*CUPS architecture*)	OLT is a standalone network element

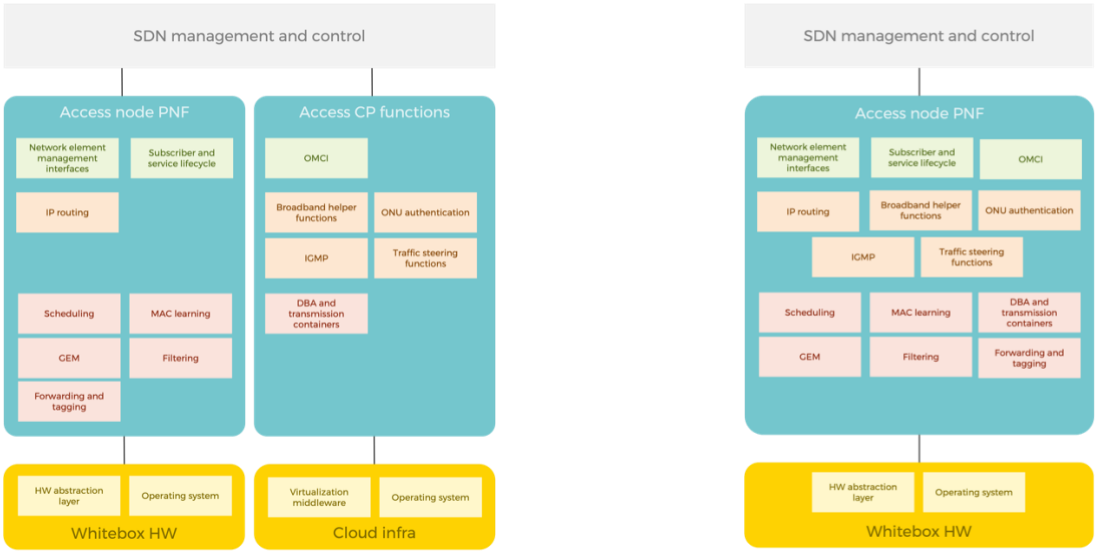

Figure 4 - Standalone vs cloudified OLT SW architectures

Per the above schematic, the OLT functions that can be located in the cloud includes Broadband Helper functions, IGMP proxy/snooping, ONT authentication, vOMCI and vDBA. Specific guidance related to CUPS implementations of these functions is provided in relevant areas of the 4 Software requirements section below. Where the CUPS architecture is deployed, the cloudified functions may be located in local servers close to the OLT, centrally in a cloud environment, or split between multiple locations.

Each operator will make a unique decision as to if, how and when to adopt a CUPS

architecture. Indeed, many operators plan to transition between these approaches during the operational lifespan of their OFAN OLTs. As such, SW solutions should provide a high degree of flexibility regarding the placement of functions, and should consider the ability to migrate and transition between different approaches over time. In particular, the software architecture of products built to address this use case must be sufficiently modular to allow individual functions to deployed on top of the NOS locally on the whitebox, and/or into a variety of different cloud infrastructure locations. Some example deployment scenarios include:

- Functions are hosted as a CNF within the whitebox itself
- Functions are hosted on an external server adjacent to the OLT
- Functions are hosted remotely in a public or private cloud

To this end, the use of open and standards-based interfaces between functional components is extremely important – regardless of whether they are deployed in a CUPS architecture, or locally on the OLT itself.

The initial intention is to enable architectural transformation, and to increase the velocity and agility of software developments in the OLT space. In the longer term, the OFAN group is interested to understand opportunities to 'mix and match' different software functions provided by different parties, and would welcome feedback and discussion with industry players on this topic.

The cloudification of OLT functions is an area of active development within the Broadband Forum, and the approach is being defined under the working standard [WT-477]. This standard is expected to define the CUPS interface, and this document will be updated to provide more specific guidance in this area once the standard has been finalized and published.

2.2.3 Security considerations

Network security is a highly important topic for any operator, and technologies deployed in a CSP network are expected to be secure by design to protect the integrity, security and availability of the services they support. In practice, this means:

- The forwarding plane must implement mechanisms to control the flow of traffic between users
- The control plane must implement mechanisms to protect the OLT from denial

of service attacks and unauthorized access, whether intentionally malicious or not

- The management plane must implement mechanisms to prevent unauthorized access to the device, and to prevent interception of management traffic

More generally, any operating system, middleware, database, application or other software used in the OLT shall be hardened according to the benchmarks provided by CIS level 2 (Center for Internet Security, http://www.cisecurity.org) Level 2 at 95% (or equivalent benchmark).

2.3 Migration, coexistence, and enhancements

A key principle of this initiative is that it should be easily deployable into both greenfield and brownfield network environments. Accordingly, we expect that HW and SW products will need to coexist alongside existing OLTs and other network components. This may be as part of a migration scenario (in the event OFAN OLTs are replacing legacy platforms), or as a more permanent coexistence strategy for solution / vendor diversity.

The OFAN requirements anticipate such scenarios by relying on open and standards-based interfaces to guarantee device-to-device and device-to-system interoperability. Indeed, it is the availability of these interfaces (defined by BBF, ONF, and others) that make this initiative possible in the first place.

While we do acknowledge that there will inherently be some differences between 'traditional' OLTs and products developed according to the OFAN requirements, we generally expect these differences won't result in impaired functionality or performance. On the contrary, we hope that by providing clarity about which features and capabilities are most important, we will empower the ecosystem to optimize and improve OLT technologies to be better and more highly performant than those that have come before.

3

Physical Requirements

3 Physical requirements

This section describes the physical characteristics of the OLT.

3.1 Form factor

The OLT will be installed within a standard 19" rack. In the case of central office deployments, 300mm/450mm/600mm ETSI standard racks are common, while remote cabinet deployments are typically installed in a 300mm deep rack. Therefore, OLTs are required to have a maximum depth of 300mm, and ideally no more than 250mm to ensure optimal airflow and cabling.

In the case of deployment into a 600mm deep footprint, OLTs are frequently installed in a 'back to back' configuration with other passive elements. This drives specific airflow constraints/requirements (see below), and means that all field-serviceable components must be accessible via the front panel to ease replacement in constrained physical environments.

The Exchange-Based OLT is a fixed form factor device, with no requirements for line card or control card modularity. Exchange and cabinet environments are typically space-constrained, so the height of the device should be optimized according to the number of ports it can provide. The 3.6 System Configurations section below provides guidance regarding port configurations that are specifically desirable, and as a general rule each device should provide a minimum of 16 ports of PON per 1RU (e.g. 16 ports in 1RU, 32 ports in 2RU).

The airflow shall be either front-to-back or side-to-side, with side-to-side deployments required where OLTs are to be deployed in back-to-back configurations. The ability to select the most appropriate fan tray(s) at the point of commissioning would be valued to simplify logistics, and as such a variety of different SKUs (front-to-back, left-to-right and right-to-left) are required.

The equipment shall conform to or exceed ETSI standard ETS 300 019-1-3 Class 3.4 for both central office and remote cabinet deployments. The desired temperature range should be (-40ºC to +70ºC) with a humidity of between 10%-100%.

The equipment shall also conform to all applicable standards regarding mechanical, electrical and safety conditions, and must comply with all directives and certifications required in the countries where it will be deployed. In particular, solutions that improve the energy efficiency of the node will be extremely valued, and this topic is covered in more detail later in the Special considerations for deployment of more power-efficient OLTs section below.

3.2 Power supply and cooling

The OLT hardware platform must include redundant power supplies and cooling components, including fans and replaceable filters to capture airborne dust and particulate matter from air intake and exhaust vents. It must be possible to substitute any of these with the device still in operation (hot swapping). Likewise, the system must be capable of maintaining full-service operation in the event of a single power feed failure. The device must generate appropriate alarms in the event of any state changes (e.g. faults, failures) within the power and cooling systems.

The equipment must support DC power supplies as the most common mode of powering OLTs in the field. The availability of optional AC power supplies would be valued for lab and AC-powered cabinet environments.

The equipment must also include variable speed blowers, to adapt the cooling capacity to the exact demand, and to optimize for power consumption.

More generally, we are also interested in approaches that can optimize power consumption - e.g. when ports are not in use, on the basis of traffic load, functional/chipset capabilities in use, etc.

3.3 Hardware platform management

The OLT hardware must include, as a minimum, one RJ45 console port, one Ethernet management port and one USB 3.0 port, which will be used for local configuration and debugging. For security reasons, it must be possible to remotely disable the console and/or the management port and/or the USB port. It is therefore mandatory that the platform also supports in-band management; in-band management should be

supported as a shared function over traffic ports, and also via a dedicated Ethernet port (SFP or RJ45).

The hardware platform must also include status indicators, including per port LEDs.

The proposed solution must implement a function that is in charge of pOLT Hardware recognition and management, and which is capable of controlling all components of the whitebox (e.g. fans, power, storage, interfaces, etc.). This function MUST be upgradeable independently from any other software module (inside or outside the pOLT).

3.4 CPU, memory

At the time of writing, we understand that 64-bit x86 is the CPU architecture that has the most robust developer ecosystem and the most appealing roadmap to ensure platform longevity. However, we are interested in exploring alternatives in this area if they are available, particularly if they can help improve power efficiency of OLT devices.

3.5 Traffic Ports

The following sections describe requirements for PON and Ethernet ports carrying service/customer traffic.

3.5.1 Ethernet Ports

The OLT provides a number of Ethernet ports. The primary purpose of these ports is to provide connectivity towards other network elements (NNI), although the ability to use some ports for connecting Ethernet-based P2P customer services (UNI) is valued.

Ethernet ports shall be provided as cages ready to accept pluggable modules. The number of cages varies according to the system configurations described below, and will consist of a combination of:

- SFP28 ports supporting 10GbE and 25GbE operation. Support for 1GbE operation is valued
- QSFP28 ports supporting 100GbE and 40GbE operation, with support for 4x25GbE and/or 4x10GbE via breakout cables.

There are a variety of possible deployment scenarios here; sometimes operators will deploy 10GE, sometimes 100GE, or sometimes both. We would invite suppliers to highlight any possible constraints or conditions that may be relevant here.

Ethernet ports shall be able to support pluggable modules that can operate at the temperature range of -40C to +70C range and configurable by software. The platform shall be compatible with a pay as you grow model defined by software licenses.

The platform shall be fully interoperable with any 3rd-party pluggable optics without restriction.

3.5.2 PON ports

The OLT provides a number of PON ports. The purpose of these ports is for connecting downstream customer services.

All PON ports must support 'flexible PON' mode of operation. This means the device must support GPON, XGS-PON or combo-PON (GPON+XGS-PON) transceivers without limitation. Support for Active Ethernet pluggables would be valued to deliver P2P Ethernet services for customers.

PON ports shall be able to support pluggable modules that can operate at the temperature range of -40C to +70C range and configurable by software. The platform shall be compatible with a pay as you grow model defined by software licenses, and the platform shall be fully interoperable with 3rd-party pluggable optics without restriction.

3.6 System Configurations

We envisage the following standard configurations are required:

	Physical size	SFP28 ports	QSFP28 ports	PON ports	
SC-1	Up to 1RU	4	2	16 – 24	Target for singleton deployments into small locations, and for more granular scaling

SC-2	Up to 2RU	4	2	32 – 48	Target for deployment as part of a clustered solution

Support for configurations offering greater port density (especially larger numbers of QSFP28 ports) would be values, and of course suppliers are welcome to propose alternative configurations based on the principles outlined in this document. We are interested to hear feedback about how these configurations can be further optimized to make best use of HW and SW capabilities.

3.7 Frequency and Time Sync Distribution

The OLT may be deployed in different scenarios requiring local and/or remote synchronization signals via local or NNI interfaces respectively. Where timing is required, the OLT shall be able to provide frequency and time synchronization to 2G/3G/4G/5G base stations which are connected to it. Suppliers are asked to consider how this functionality could be made optional - either via licenses, or via specific SKUs. The platform shall support the following time sync requirements:

- Network quality model (microsecond precision) according to [G.8271.1] Network conditions and reference model where the Boundary Clock is requested to work.
- Node performance (noise generation, tolerance, transfer and holdover) according to [G.8273.2] Section 7.1/7.2/7.3/7.4.
- Node performance (upon wander, failure, holdover) according to [G.8273.2] (7.2/7.3/annex) Boundary clock quality objectives in holdover mode
- Interoperability based on IEEE 1588 profile defined in [G.8275.1] with Boundary Clock and SyncE support for holdover purposes and Grandmaster redundant sources support. This requirement corresponds to the support of IEEE 1588v2 profile for telecoms – Precision Time Protocol and includes SyncE support in Ethernet interfaces as mandatory as per [G.8262] and [G.8264].

The OLT shall provide a 1PPS in/out external sync interfaces, and shall also support a SFP input for GPS. The platform shall be interoperable with most of the SFP providers' solutions in the market in order to avoid interoperability issues. This GPS input shall be used only in scenarios where the time sync signal cannot be received from the backhaul network through a standard Ethernet traffic port in case time sync distribution mechanisms are used.

The ability to support time distribution over GPON and XGS-PON is valued, based on G.984.3, G.9807.1 and G.988. In this case the OLT performs the slave port function, synchronizes the PON line rate to the network clock frequency and transfers the time of day information to the ONT using the method in G.9807.1 clause C.13.2.

3.8 Forwarding capacity

The OLT shall be able to cope with current and future traffic demands for fixed line networks. In practice this means the forwarding engine shall provide sufficient capacity to support full duplex forwarding at line rate between PON ports and Ethernet ports without any limitations. This includes appropriate buffer / packet memory to support forwarding between interfaces of differing rates (e.g. 100GE Ethernet to lower speed interfaces, e.g. GPON).

3.9 Special considerations for deployment of more power-efficient OLTs

It should be noted that OLTs may be installed within buildings that have limited forced-air cooling/air-conditioning capability. While such locations represent a hostile environment for the hosting of telecoms equipment, some of these locations may be less constrained in terms of floor space, which would allow for OLTs to be deployed in lower density configurations to allow for limited power/cooling capacity per footprint. By extension, we are interested in exploring options for OLTs with lower density port counts (both NNI and PON) where such units will dissipate lower waste energy to the environment as a function of RU rack-space when compared to more compact OLT designs.

It is envisaged that such low density OLTs should operate at ambient temperatures commonly found in these environments (typically up to 25C) without the need for any environmental conditioning - especially if traditional passive heat dissipation systems are used (such as heat deflector plates). It is also expected that such power-efficient OLTs will not contain fans at all i.e. be entirely passive. Due to the lower port densities, it is expected that heatsinks could be used on say the front and back of the chassis to passively dissipate the internal heat energy.

It is recognized that such passive low density OLT designs may be more expensive to

build, particularly considering the need for higher temperature specified components and the larger chassis with substantial extra metalwork in the form of heatsinks/heatpipes. However, we believe that the higher initial capital costs could be offset by lower OPEX costs (both from electricity usage and air-conditioning requirements).

Additionally, the expected operating lifetime of the next-generation OLT is probably going to be considerably shorter than has been the case for traditional telecoms access network equipment. While the lifespan of a traditional OLT is expected to be 7-10 years (or more) we believe that in some deployment scenarios a shorter refresh cycle may be acceptable for such a device, particularly considering constantly increasing bandwidth requirements and the rapid evolution of new PON technologies.

Note that the requirements for the power efficient passive OLT are not the same as for remote mounted OLTs (that are expected to operate from within a sealed outdoor plant chassis) since it is expected that such OLTs will be operating from open racks within manned buildings and co-located with other telecoms network equipment. However, it is expected that the passive nature of the design could take some elements from the design of such entirely sealed units.

While the OLT is only a part of the problem, we believe that such temperature-hardened OLTs could be an interesting solution to address evolving conditions in both central office and remote cabinet environments. We would value further discussion, feedback, contributions, and ultimately commercial products from suppliers that could address such requirements.

4

Software Requirements

4 Software requirements

This section describes the SW requirements for OLT systems realized according to the different possible SW architectures described above. For the purposes of this document, we split the SW into three main sections:

- **User Plane**, which describe the requirements for Ethernet/PON forwarding plane functions and PON access line functions
- **Control Plane**, which describes requirements for Layer 2 (Ethernet) and Layer 3 control plane functions
- **Management Plane**, which describes requirements for the OMCI interface, as well as broader management of the OLT itself

User plane / Forwarding plane	Forwarding and tagging	Physical and logical topologies in which the OLT is deployed
	MAC learning	Requirements for MAC bridging and learning
	Filtering	Access control lists and other frame / packet filters
	Scheduling	Quality of service requirements
User plane / Access line	DBA and Transmission Containers	Dynamic Bandwidth Allocation
	GEM	GEM port mappings
Control plane	Broadband helper functions	PPP-IA and DHCP-RA requirements
	IGMP	Ethernet multicast requirements
	IP routing	IP routing protocols
	ONT authentication	ONT identification and bootstrap
	Traffic steering	WT-474 traffic steering
Management plane	OMCI	ONT management
	Network element management	SDN, alarms and performance management
	Subscriber lifecycle	Service provisioning and assurance

4.1 User Plane requirements

This section describes the functional and non-functional requirements related to user plane functions. These requirements apply equally to the centralized and distributed deployment models.

4.1.2 Forwarding and Tagging

The following section describes how the OLT interacts with traffic, including functional and non-functional (e.g. scalability) requirements.

4.1.2.1 Ethernet topologies and bundling

The OLT integrates with the transport network via Ethernet connections. A variety of different scenarios are possible, depending on the underlying topology of the operator's backhaul and transport networks. This section of the document will explore some of the most common topologies, and the associated requirements for the OLT.

Each OLT may be connected to just a single upstream transport node, or may be connected to multiple adjacent nodes, with a mixture of upstream- and downstream-facing connections. Depending on the upstream network architecture, these upstream devices may operate in an active/standby configuration, or as an active/active cluster.

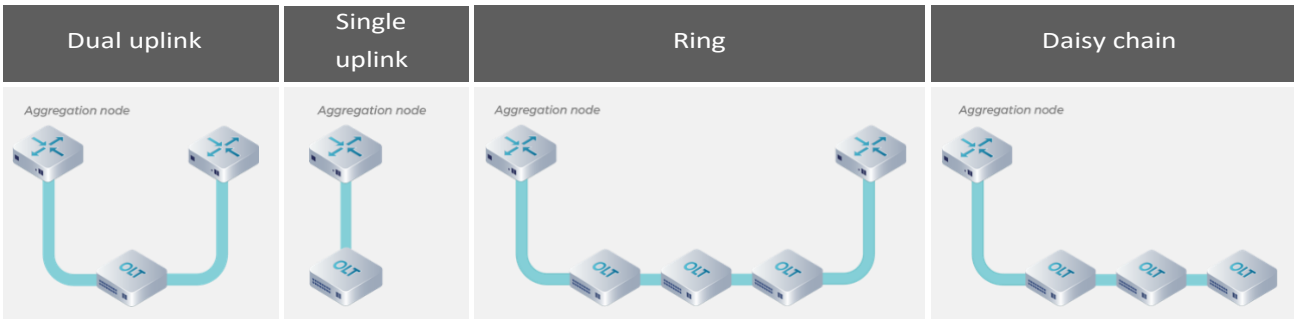

Figure 5 - OLT uplink and subtending topologies

Some indicative scenarios are represented in the schematic above. However, more complex hybrid scenarios are also possible and the OLT should support a wide range of possible deployment architectures without limitation.

A given OLT NNI link may be provided as an aggregated/bundled configuration for capacity/resiliency reasons, or for other operational reasons.

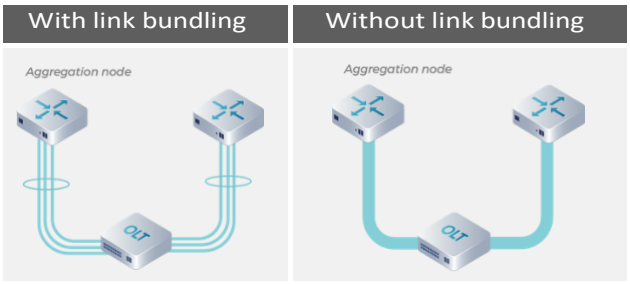

Figure 6 - Candidate link bundling topologies

In these cases, the bundling may be based on LACP or static configuration. In the case of LACP, the OLT shall support both active/active and active/standby modes of operation, including the ability to integrate with remote multi-chassis LAG implementations. The ability to construct LAGs consisting of different interface variants (e.g. 10GbE+100GbE) is valued.

The OLT shall support jumbo frames with MTU up to 9216 bytes.

4.1.2.2 VLANs and bridging

The Ethernet connection between the OLT and the transport network uses VLAN tags to differentiate between customer, service and traffic types. Accordingly, the OLT must support encapsulating Ethernet frames according to the IEEE 802.1Q standard to enable the connectivity with other network elements using multiple virtual LANs.

The OLT must support multiple VLAN models for subscriber attachment:

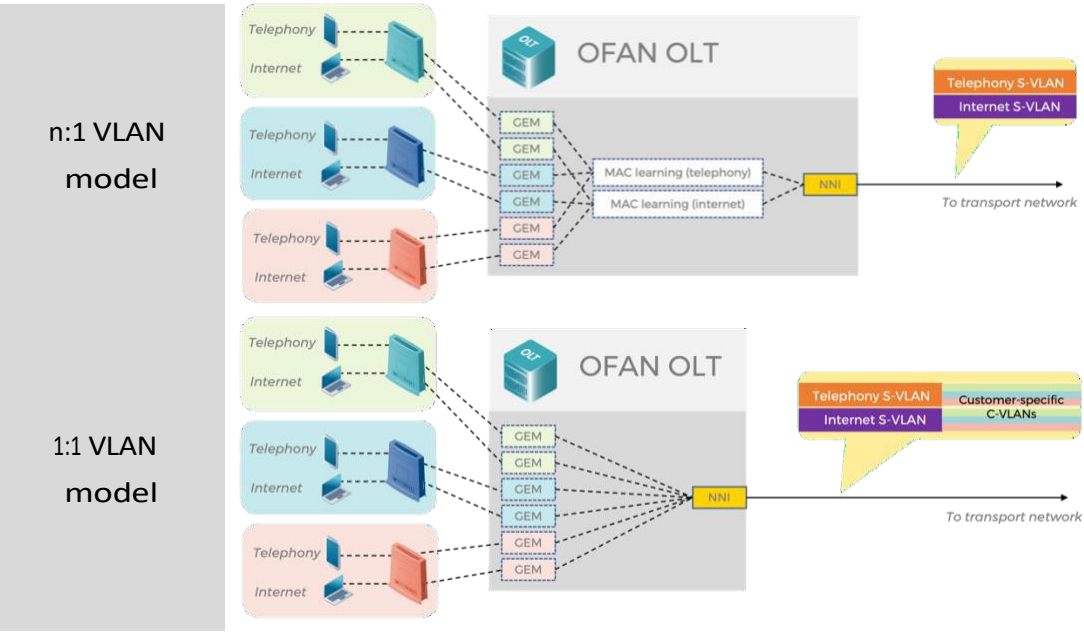

Figure 7 - VLAN attachment models

In the n:1 model, multiple subscribers are mapped to a common S-VLAN, typically using a unique S-VLAN for each service type. This may be used for delivering simple residential services, or as part of a layer 3 / IP routed service architecture. In the 1:1 model, each subscriber service is mapped into a unique S+C-VLAN. Services carried using S+C-VLAN tags require support for VLAN stacking (according to IEEE 802.1ad and its evolutions), and support for MAC address hiding/aggregation is desirable to limit the number of MAC addresses the OLT presents to the upstream transport network.

This VLAN mapping is configured at the granularity of a specific GEM port. Where a given subscriber has multiple services, it is common for each service to be carried using a different VLAN model - e.g. a unicast/data VLAN may use 1:1 VLAN, while multicast traffic may be carried using a shared n:1 VLAN.

The OLT must support a variety of push/pop operations for both S-VLAN and C-VLAN tags on NNI ports, and on any directly connected UNIs (e.g. for point-to-point Ethernet services). This includes support for tagged and untagged interfaces. Support for VLAN swap operations is valued.

4.1.2.3 Multicast

IP multicast is widely deployed for the delivery of linear IPTV services, and the OLT must be able to receive IP multicast streams on every network interface and to replicate them to GPON/XGS-PON ports towards the ONTs.

There are various multicast forwarding schemes that can be deployed, as depicted below.

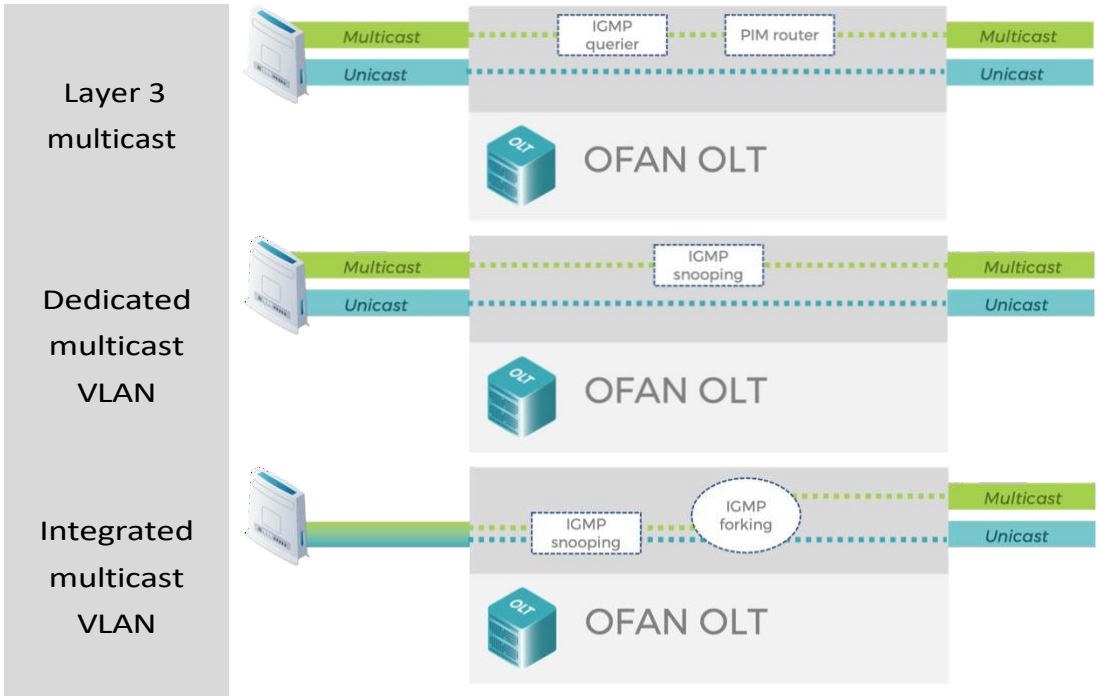

Figure 8 - Multicast forwarding models

The first scenario is where multicast traffic is carried at layer 3, with IGMP query and PIM capabilities provided by the OLT. This scenario is described in more detail in the 4.2.4.4 IP multicast section below.

The next scenario is the dedicated multicast VLAN model, where multicast traffic for multiple subscribers is carried in a unique VLAN that is dedicated for multicast traffic. This topology is common where customer multicast devices are carried in a unique VLAN in the home, or where IGMP forking is enabled in the CPE/RG. This scenario requires that the OLT implements IGMP snooping (see 4.2.3 IGMP below) to identify which RGs (i.e. ONTs) have subscribed to a given multicast group.

The final scenario is the Integrated multicast VLAN model, where multicast traffic shares the same VLAN as unicast traffic. Again, IGMP snooping is required in the case of subscriber services attached using the n:1 model. In this scenario, the OLT may be required to support multicast forking to enable separation of multicast traffic within the upstream transport network.

4.1.3 MAC Learning

The most typical design for an OLT is to configure a single bridge domain, which implements a common MAC address table and VLAN tag space for NNIs upstream towards the transport network and/or downstream towards subtended OLTs.

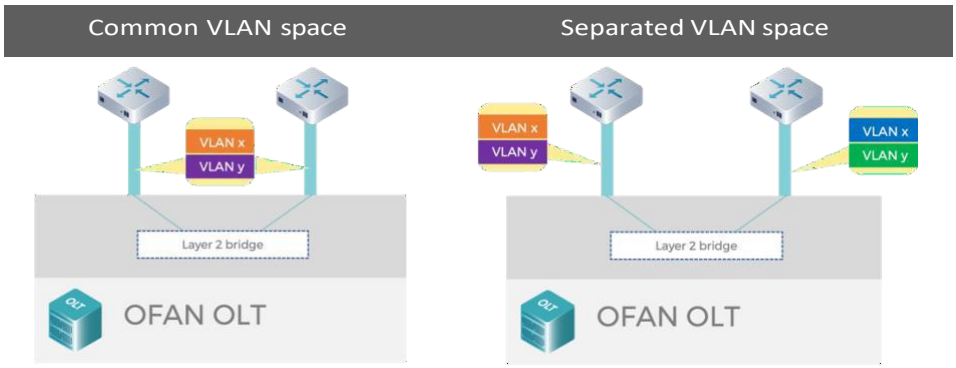

Figure 9 - Layer 2 bridge / MAC learning models

The ability to support multiple, separated bridge domains within a single OLT is valued - e.g. to support Virtual Unbundled Local Access (VULA) wholesale scenarios. These separated bridge domains allow for VLAN tags to be re-used across multiple links, and for each VLAN + link combination to be implemented as a unique layer 2 bridge.

4.1.4 Scheduling

Fixed operators rely on quality of service technologies and architectures to implement differentiation between different application, service and customer types.

Classification	Scheduling	Marking
 *Identification of traffic types based on their characteristics*	 *Controlled and differentiated transmission of data*	 *Signaling QoS behaviors to remote network elements*

Figure 10 - OLT QoS capabilities

The QoS architecture of the OLT implements three main functions:

The first function is **classification**, whereby the device determines what type of traffic it is seeing. For this use case, the OLT must support classification based on:

- 802.1P values contained in the VLAN headers in the downstream direction (i.e. received from an NNI). This is known as 'trusted mode'. The ability to flexibly configure the OLT to use the value from either S-VLAN or C-VLAN tags (where used) is valued
- The ingress port/link over which the traffic was received, including the ability to classify traffic based on logical link identifiers such as VLAN ID. This is known as 'untrusted mode'.

Support for other forms of classifier (e.g. DSCP, IP ToS) would be valued. OLT must support trusted (preserve tags) and untrusted (overwrite tags)

The second function is **scheduling**, whereby the device implements specific queueing behaviors for each traffic class. The OLT must support, at a minimum, strict priority (SP) scheduling and some kind of weighted queue scheduler (weighted round robin - WRR, or weighted fair queueing - WFQ), along with the ability to apply shapers and policers. Support for dual-mode (SP+WRR/WFQ) schedulers and other scheduling modes is highly valued.

Given the increasingly diverse nature of fixed line services, the OLT must support hierarchical scheduling, with the ability to combine multiple schedulers. Common use cases for such a hierarchical QoS architecture include:

- The need to combine multiple scheduling modes in a single subscriber context (e.g. a WRR to differentiate Internet vs multicast data traffic, and a PriQ to differentiate aggregated data traffic vs voice traffic)
- The need to combine traffic from a group of subscribers under a common scheduler (e.g. to differentiate between traffic from business subscriber services,

and residential subscriber services)

The final function is **marking**, whereby the device marks frames and packets with QoS tags that will be used by onward network elements. As with classification, the primary mechanism here is 802.1P, and the OLT must be able to remark 802.1P bits in both upstream and downstream directions, including:

- The capability for trusted (preserve tags) and untrusted (overwrite tags) mode of operation.
- The ability to set the value of 802.1P values in both C-VLAN and S-VLAN headers (where used). This includes the ability to set each header value independently of the other, and the ability to map/reflect from one to the other

As with classification actions, the ability to interact with other QoS and ToS bit identifiers is valued.

The OLT must support all of the above capabilities in both upstream and downstream directions, with each direction being configured differently where needed.

4.1.5 Filtering and Access Control List (ACLs)

Packet filtering plays a vital role in helping operators secure their infrastructure and their customers. These technologies also form part of how operators implement and control their high order service offerings - a common use case is to control access to multicast streams for IPTV services.

Security use cases typically require ACLs that are common across multiple/all subscribers; these filters may be applied on any interface - NNI, subscriber port, or even towards the control plane. Meanwhile, filters used to support service offerings are applied to a subscriber port only, but may require a unique ACL to be applied for each user, or group of users. As such, it's important that the OLT can support a granular and flexible approach to defining and applying packet filters.

As a baseline, the OLT must allow filters to be specified in the upstream direction (from customers), and/or the downstream direction (towards customers). A variety of different packet filter types is needed, including:

- User defined L2 filter, including src MAC address (+prefix), dst MAC address (+prefix), ethertype, p-bit, VLAN ID.

- User-defined L3 filter including src IP Address (+prefix), dst IP address (+prefix).
- User-defined L4 filter including src/dst IP address, L4 protocol and src/dst ports.

In all cases, it should be possible to extract counters and other diagnostics/operational data from the OLT, that show which ACL entries / lines have been matched.

4.1.6 Scalability and performance

The following section describes the minimum targets for scalability of forwarding plane functions.

4.1.6.1 Maximum delay

The maximum round-trip delay introduced for a service must not exceed 2ms. This RTT is measured between the NNI port on the OLT and the UNI port on the ONT, and is related to guaranteed bitrate services only (i.e. exclusive of congestion caused by oversubscription). We are interested in exploring opportunities to reduce or otherwise optimize this delay to increase QoE for customer services and/or to support low-latency applications.

4.1.6.2 Forwarding plane scalability

The forwarding plane shall be scalable according to the number of subscribers and services connected to each OLT device. With reference to the standard configurations provided earlier, this means:

- Support a high number of MAC addresses. The reference number is 32k MAC addresses for an OLT with 16 ports.
- Support 4096 VLAN identifiers in user, service and stacked VLANs.
- Support 1024 multicast groups.

4.1.7 Access Line functions

4.1.7.1 General characteristics

The optical interface of the OLT must be according to the GPON [G.984.x] and XGS-PON [G.9807.1] standards.

The OLT must support Flexible PON technologies, supporting GPON, XGS-PON or Multi-PON modules (MPM) in a flexible way at a PON port basis. The OLT modules will support GPON optical classes B+, C+, D as specified in [G.984.2], XGS-PON optical classes N1, N2 as specified in [G.9807.1] and Multi-PON optical classes B+, C+ as specified in [G984].

The OLT must support a wide range of split ratios on each PON port, up to a minimum of 1:128, and including support for resilience and protection of the ODN (Type B and optionally Type C). Additionally, the OLT must support Forward Error Correction (FEC) in downstream and upstream direction, including the ability to configure each independently of the other.

4.1.7.2 DBA and Transmission Containers (T-CONT)

The OLT shall support a minimum of 1024 T-CONTs per port in GPON and 2048 in XGS-PON, with 8 priority bits per T-CONT. This shall include support for status reporting (SR) and traffic monitoring (TM) to manage the dynamic activity of upstream traffic from ONTs.

The OLT shall also support DBA for sharing the upstream bandwidth among the connected ONTs, providing strict-QoS that enables control of upstream capacity, latency and jitter for each T-CONT.

In the case of a CUPS architecture, the DBA function could be disaggregated into a vDBA algorithm that runs as a remote software component, and an engine partially located in the pOLT, which is assumed to mainly process data at data plane level (as described in the Figure 8-1 of [TR-402]). In this scenario, the vDBA algorithm provides strict-QoS, while the physical OLT takes care of framing and other data plane functions.

The virtual DBA (vDBA) algorithm controls the behavior of the engine by computing a virtual bandwidth map (vBMap) for each slice of the pOLT and delivers the vBMap to the engine on the pOLT. Multiple vDBA functions could operate with the same pOLT. The vBMap indicates the desired position for each slot allocation within a slice, while the PHY-BMap is the allowed upstream capacity in the pOLT.

The implementation of a vDBA algorithm is according to [TR-402] and [TR-403] and further amendments, and section 4.5.5 of [WT-477].

4.1.7.3 GEM

The OLT shall support a minimum of 4000 GEM ports per PON port, including the ability to perform GEM port mapping based on VLAN, 802.1p-bit and VLAN+p-bit. This also includes support for multicast GEM ports

4.1.8 User Plane Security requirements

The OLT shall allow the possibility to prevent the direct connectivity of the users at PON and OLT level, ensuring that visibility among PON users is only possible at the L3 edge (e.g. direct ping and remote access will not work between users of the same OLT).
As the downstream traffic arrives at all the ONTs connected to the same PON port, any user can intercept the traffic intended to other ONTs. In order to guarantee the confidentiality of the information, the OLT shall support AES-128 bits encryption in the downstream channel in GPON and in the downstream and upstream channels in XGS-PON. In general, the use of encryption is configurable depending on operator requirements.

4.2 Control Plane requirements

When describing the control plane requirements for a disaggregated OLT, the reference architecture described in section 2 Platform Architecture applies.

4.2.1 Broadband Helper Functions

4.2.1.1 DHCP Relay Agent

According to [RFC3046] it must be possible to enable/disable the DHCP Relay Agent functionality, per OLT GPON/XGSPON port, GEM port or VLAN (SVLAN) basis, without any limitation on the number of VLANs, GEM ports or GPON/XGS-PON ports.

The DHCP RA function inserts the “option 82” on all the messages sent by the DHCP

client located at the customer side, typically using “Agent Circuit-ID” sub-option 1 and/or “Agent Remote-ID” sub-option 2. Likewise, the DHCP RA function removes all the “option-82” information from all DHCP reply messages received from the DHCP server before forwarding them to the client.
According to [RFC6221], it must be possible to enable/disable the Lightweight DHCPv6 Relay Agent (LDRA) functionality, without any limitation on the number of OLT VLANs, GEM ports and GPON/XGS-PON ports. Typically the LDRA function inserts DHCPv6 Interface ID (option 18) or DHCPv6 Remote ID (option 37) to the messages sent by the client to DHCP server.

The required behaviors are defined in [TR-101] and [TR-156]. In the case of a CUPS architecture, DHCP relay agent functions may be implemented as a module within a cloud environment, and are invoked per section 4.5.1 of [WT-477].

4.2.1.2 PPPoE Intermediate Agent

As per [TR-101], it must be possible to configure a PPPOE IA, on a per OLT VLAN (SVLAN) basis, in order to allow the access node to insert access loop identification within the PPPOE frames during the protocol discovery stage.

In the case of a CUPS architecture, DHCP relay agent functions may be implemented as a module within a cloud environment, and are invoked per section 4.5.2 of [WT-477].

4.2.2 Ethernet Link Protection

According to IEEE 802.1d, IEEE 802.1w and IEEE 802.1s standards, it must be possible to turn on/off Spanning-Tree protocols such as Rapid STP (RSTP) and Multiple STP (MSTP) for L2 link protection on uplinks. Also, support for [G.8031] Ethernet Linear Protection Switching is required. Support for Ethernet Ring Protection Switching (ERPS), as per [G.8032] recommendation, is valued.

4.2.3 IGMP

In order to enable an efficient delivery of multicast channels to ONTs, it is requested the support of IGMPv2 [RFC2236], IGMPv3 [RFC3376] and MLDv2 [RFC3810], augmented with snooping and proxy capabilities.

As such multicast protocols are vulnerable to attacks, any specific measures to prevent security flaws are valued.

In order to ensure proper multicast service quality and availability it must be possible to configure a multicast Call Admission Control (CAC) functionality, e.g. based on the max number of multicast group memberships.

In the case of a CUPS architecture, IGMP proxy / snooping functions may be implemented as a module within a cloud environment, and are invoked per section 4.5.3 of [WT-477].

4.2.4 IP Routing

Current solutions of open and disaggregated OLT support layer 2 functionalities; however, there are some use cases where operators may prefer to use a layer 3 service architecture, for example for IPTV and/or VoIP services. In those cases, an OLT supporting the following layer 3 functionalities would be preferable as it would be possible to maintain the same end to end network and service architecture for both open OLTs and legacy OLTs. In that case, the migration can be done smoothly and the coexistence may be easier without affecting other network layers and without a service redefinition because of the introduction of open OLTs.

The following sections describe the critical considerations for a layer 3 OLT. Support for some or all of these functions would be highly valued as a means to improve service flexibility. In these cases, the OLT should support IP unicast packet routing from the UNI to NNI and vice versa, providing layer 3 routing with MAC transparency and the ability to forward traffic for any VLAN. Note that the RIB/FIB requirements for such a device are expected to be relatively modest (e.g. 16k IP routes for an OLT with 16 PON ports).

In the case of CUPS deployments, IP routing functions (primarily routing and MPLS protocols) functions could potentially be virtualized. At present there isn't any specification that defines the CUPS protocols for IP routing functions in OLTs, although this will be most likely done in one of the future releases of [WT-477]. In the meantime,

we welcome any inputs from the community on how to virtualize layer 3 and MPLS functions.

4.2.4.1 Routing Protocols

The OLT should support static and dynamic routing. The OLT should interchange routing information and perform path calculation tasks.
The OLT should support at least 16k ARP (Address Resolution Protocol) entries (reference number for an OLT with 16 PON ports).
The OLT should support the following routing protocols and L3 functionalities:

- IGPs such as RIPv6, OSPF and IS-IS. Since an OLT is typically deployed as part of a regional metro network segment, we would expect to see a minimum of 100 adjacencies.
- BGPv4, with support for a minimum of 50 BGP sessions.
- VRF (virtual routing and forwarding) for the purposes of traffic isolation, and RSVP for traffic engineering purposes
- BFD for rapid failure detection

4.2.4.1.1 MPLS

MPLS is a switching protocol that operators may use in OLTs to provide certain services or to provide new services in the future. An OLT supporting MPLS should fulfil the following requirements:

- Support for LDP and RSVP for the signaling and establishment of MPLS tunnels
- Support for layer 2 VPN services delivered using VPLS and pseudowires
- Support for MPLS (tagged or untagged) and VLAN-tagged traffic on the same physical port
- Support for EXP as part of QoS classification and marking functions described in section Scheduling

While this functionality can typically be provided by an external top of rack switch aggregating traffic from multiple OLTs, we are interested in opportunities to incorporate MPLS functionality into the OLT in the most cost-effective fashion, and the pros and cons of the different approaches. We would welcome feedback about how this could be achieved - e.g. via virtualization of SW functions, creation of more powerful HW SKUs, or other alternatives.

We will provide more information about specific MPLS features in the detailed technical requirements document.

4.2.4.2 IP filtering

The OLT should support IP filtering with at least 64 filtering rules per VLAN (reference number for an OLT with 16 PON ports).

4.2.4.3 QoS

The OLT should support simultaneous translation of CoS and DSCP.

4.2.4.4 IP multicast

The OLT should support IGMP version 2 and version 3 multicast protocols in accordance with [RFC2236] and [RFC3376] respectively.

The OLT should support a mechanism to immediately resign from the multicast groups ("IGMP immediate-leave") and to allow the subscription to the groups from IP addresses not directly connected ("IGMP promiscuous"). The OLT should be able to replicate IP multicast packets towards the ONTs subscribed to each multicast group in multiple PON ports. The same IP multicast channels may be repeated in different VLANs at the same time.

In addition, the OLT should be able to learn all the required information about the multicast groups available at any time including the multicast IP address and its source IP address.

The OLT should support a mechanism to grant and deny potential multicast sources by source IP address.

The OLT should support the following modes of PIM (Protocol Independent Multicast):

- PIM-SM (Sparse Mode) in accordance with [RFC7761].
- PIM-SSM (Source Specific Multicast) in accordance with [RFC4607].
- PIM-ASM (Any Source Multicast) in accordance with [RFC8815].

4.2.5 ONT Authentication

ONT authentication ([G.984.3], [G.987.3] and [G.988]) must be supported using serial number, PLOAM password (Registration ID for XGS-PON) and serial number plus PLOAM password. The authentication method shall be configurable per ONT.

In the case of CUPS deployments, [TR-489] describes multiple scenarios for ONT authentication. The virtualized solution only implements scenario 4 as described in section 4.5.4 of [WT-477].

4.2.6 Traffic Session Steering (TSF)

The Traffic Steering Function is primarily a user plane component, responsible for forwarding the traffic related to a specific subscriber context between the OLT where the subscriber is attached and the specific service gateway (e.g. BNG) user plane that has been identified through the steering process.

Traffic Steering architecture and interfaces are under development in [WT-474]. While this work is still ongoing, there are some foundational elements that are of note for this requirements document:

- The ability to create a variety of different types/variants of network connections between the OLT and the service gateway
- Granular and flexible means to classify subscriber contexts
- Implementation of appropriate control interfaces

Early implementations of these features is valued, ahead of more complete implementations as the work to develop the [WT-474] and [WT-477] standards progresses in the Broadband Forum.

4.2.7 Control Plane security requirements

To protect the control plane of the OLT, the device must Include appropriate measures to protect against IGMP, ICMP and DoS attacks and MAC spoofing in protocols (ARP, IGMP, DHCP, etc.). This includes support for flexible control plane policing / protection policies, which can be configured to limit and/or filter the forwarding of packets from traffic ports towards the control plane / CPU.

4.3 Management Plane requirements

When describing the management plane requirements for a disaggregated OLT, the reference architecture described in section 2 Platform Architecture applies.

4.3.1 M&C System Specifications

The OLT is typically managed over its northbound interfaces using a management plane controller system. This system may be a traditional graphical-style EMS, or an SDN controller leveraging APIs provided by the network. The OLT may also be accessed directly from dedicated applications using CLI or API.

In any case, we expect the OLT to expose the same functionalities, regardless of what northbound systems are driving those interfaces. We also expect that the OLT should implement all management functions described without impact to service operations.

As a general principle, we expect that these requirements should be agnostic to any specific management plane system or architecture. Likewise, the same set of requirements should apply to all types of OLT deployment, including both control plane and user plane components of a CUPS architecture.

4.3.2 Northbound Interface requirements

The northbound interfaces for the solution are based on [WT-413i2], [TR-383], [TR-385] and [TR-454]. These interfaces are described in the 2.2 SW functions and reference interfaces section, including the applicable compliance references.

In general, all management interfaces (M2M or H2M) should be accessible in-band via NNI ports or out-of-band via local management Ethernet ports.

4.3.2.1 Command Line Access

In addition to programmable NBIs, the OLT can also be managed via a CLI, which can be accessed via a number of different mechanisms, including:

- Remote management through SSH
- Local out-of-band management through Ethernet and RS-232 console

4.3.3 Fault, Performance and Alarms Management

The OLT must support northbound interfaces that can be used to send telemetry, performance/fault data and alarms to a northbound management system. These northbound interfaces may be implemented using gNMI and/or gRPC, and the OLT should also be able to generate log messages via syslog. Support for IPFIX would be valued.

The OLT shall implement a flexible alarm handling system that allows for alarms to be tracked through their lifecycle. This includes:

- The ability to configure which alarms should be monitored / generated
- Support for alarms to be cleared once the fault condition has been recovered
- The ability to buffer and store alarms for a period of time in the case access to the management plane is lost

4.3.3.1 Device commissioning and ZTP

The device must implement the TIP ZTP greenlight solution for initial bootstrap and configuration of devices.

4.3.4 PON management

The PON system (OLT/ONT) supports:

- Monitoring of the transmitted and received signal strength indication (RSSI) of PON ports in operation mode.
- PON-ID, PON-TAG to identify the OLT inside the optical distribution network
- ONT firmware upgrade through the PON interface using OMCI (including automatic multiple upgrades per ONT vendor, model and firmware version)
- The alarms and the performance monitoring according to section 11 of [G.984.3] and [G.9807.1] annex C.14
- Traffic counters and statistics at PON and Ethernet interface levels, such as number of received frames, bytes, dropped frames, multicast frames, FEC errors, upstream BIP errors, downstream BIP error, etc.
- Rogue ONT prevention, detection, isolation and mitigation techniques

- ONT autodiscovery
- OLT backup, restore and software/firmware upgrades, with the ability to drive these activities without interrupting the operational state of the OLT

4.3.5 ONT management (OMCI)

The PON management interface will be according to [G.984.4] and [G.988], including using the remote ONT transceiver parameters regarding optical layer supervision (optical TX/RX, voltage, laser-bias current, temperature) according to Appendix IV of [G.984.2] Amd. 2 and Appendix B.II of [G.9807.1] for GPON and XGS-PON respectively. The OMCI function can be virtualized according to [TR-451], which defines multiple functionalities (e.g. vOLTMF, vOMFI Function, etc.). In this scenario, the vOMCI function is deployed in a remote cloud instance, while the vOLTMF can be a feature of the Access SDN M&C, a BAA Layer, or an external function. The vOMCI Function is here mainly responsible for:

- Receiving from the vOLTMF commands issued towards the target ONT(s).
- Translating the received management commands into OMCI management entities (ME) and formatting them into OMCI messages compliant with [G.988].
- Sending the management data and information generated from the received OMCI messages to the vOLTMF.

The provision of OLT-ONT interoperability is a major requirement for this project. This can be achieved using Open OMCI which in turn predicates the use of both an OpenONT and OpenOLT adapter as part of [VOLTHA] whilst the direct impact on the OLT is the requirement for interoperability with the installed OpenOLT agent. This is very much an ONF view of OMCI management including no use of vendor proprietary Managed Entities as detailed in the AT&T Open OMCI specification. Discussion of power shedding/low power mode support as per [G.Sup45] is welcomed.

4.3.6 Subscriber lifecycle management

The following sections describe the capabilities exposed by the OLT in support of fulfilment and assurance of customer services. These capabilities will typically be driven by the northbound management plane system (e.g. an EMS or SDN controller)

4.3.6.1 Fulfillment

The fulfillment process involves the configuration and activation of the different network resources managed by the OLT and their allocation for the services associated with each customer. The platform shall support the following requirements:

- The OLT must allow the configuration and activation of its resources via an external manager or controller using a northbound management plane interface. Additionally, OLTs must support local configuration and activation without the need of such a manager or controller.
- The OLT must have the capability to notify a northbound management plane system in the event of any configuration changes.
- The configuration and activation interface must allow individual resource activation, configuration and deactivation. Also, it should allow changing the current configuration of any resource and rollback of all individual resource operations. Typical resources are ONTs, VLANs, GEM Ports, T-CONTs, QoS policies, ACLs, etc.
- Once an operation has succeeded, changes must stay permanent locally at the OLT, even after system reboot.

4.3.6.2 Assurance

Main requirements involved with assurance include performance monitoring and fault monitoring are:

- The OLT must allow an external manager or controller to obtain metrics and counters using a northbound management plane interface. Alternatively, the OLT can be configured to push the selected metrics and counters to that manager or controller, either by sending them directly or by streaming them over a message broker. The time range to collect the data exported and the granularity of the measurements must be configurable.
- The OLT monitoring functionality shall be able to:
    - Retrieve statistics and performance monitoring parameters.
    - Retrieve inventory and environmental parameters
    - Report alarms, including related parameters (timestamp, severity, etc.)

also filtered on the basis of certain parameters
    - Determine OLT throughput at different levels.
    - Detect changes of state in the connections, interfaces, ports, etc.
    - Report the quality of the selected customer resources, such as lines, services, etc.
    - Obtain information about the hardware resources (such as parameters related to the chassis, PSUs, fans, transceivers, etc.).
    - Inform about the parameters related to TX/RX traffic at packet and byte levels, including errors at PON port and uplink port level.
    - Acquire information related to ONTs past events and current status.
    - Include parameters for GPON and XGS-PON performance monitoring according to section 11 of [G.984.3] and [G.9807.1] annex C.14 (XGS-PON).
- OLT alarms shall have a unique ID and shall include additional parameters such as an alarm name, category, status, network element associated to the alarm, severity, event type (alarm, notification, etc.), date and time and rest of information of [X.733].
- The OLT must allow an external manager or controller to obtain active alarms according to a filter (for example the alarm severity) using a northbound management plane interface. Additionally, the OLT must support its configuration to transmit alarms activation and modification to that manager or controller.
- It should be possible to configure the parameters of events which generate alarms, such as thresholds for alarm detection and their corresponding alarm priorities.
- The OLT must keep an active alarm storage with a global list of all active alarms.
- The OLT must support alarms related to GPON and XGS-PON according to section 11 of [G.984.3] and [G.9807.1] annex C.14 (XGS-PON).

4.3.7 Management Plane Security Requirements

- Centralized Authentication System / Server (CAS) support for authentication and authorization of the users accessing the OLT and generation of the accounting events of user activity. The supported CASs shall be mainly based on TACACS+ [RFC1492], RADIUS [RFC2865] or LDAP [RFC2307].
- The system shall support the configuration of an alternative authentication

method (for example, local authentication) in case the CAS is unavailable and only in that case.

- Only unbroken hash-like encryption algorithms (SHA-2, etc.) shall be used for equipment's local passwords, SSH keys, TACACS, RADIUS, LDAP, etc. shared secrets, control/data authentication protocol passwords / keys, SNMP communities or any other protocol.
- In case SNMP is used, SNMPv3 shall be used instead SNMPv1 or SNMPv2.
- Applications accessing the OLT via an API can only access the functions, URLs, services, resources and information for which they have specific access authorization. The client application shall be authorized by strong methods.
- Access security and anti-theft: in general, the solution must support the necessary security mechanisms to authenticate and encrypt communications between the network element and its management system or controller. The network element should offer the possibility of only enabling local traffic after the device has been authenticated by the management platform/controller. The system should also offer the possibility to enable anti-theft mechanisms that prevent the use of the equipment in any other environment than the one it was conceived in.

5 References

[G.8262] Timing characteristics of a synchronous equipment slave clock, ITU-T
https://www.itu.int/rec/T-REC-G.8262

[G.8264] Distribution of timing information through packet networks, ITU-T
https://www.itu.int/rec/T-REC-G.8264

G.8271.1] Network limits for time synchronization in packet networks with full timing support from the network, ITU-T
https://www.itu.int/rec/T-REC-G.8271.1

[G.8273.2] Timing characteristics of telecom boundary clocks and telecom time slave clocks for use with full timing support from the network, ITU-T
https://www.itu.int/rec/T-REC-G.8273.2

[G.8275.1] Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, ITU-T
https://www.itu.int/rec/T-REC-G.8275.1

[G.984.2] G-PON: Physical Media Dependent (PMD) layer specification, ITU-T
https://www.itu.int/rec/T-REC-G.984.2

[G.984.3] Gigabit-capable passive optical networks (G-PON), ITU-T
https://www.itu.int/rec/T-REC-G.984.3

[G.987.3] 10-Gigabit-capable passive optical networks (XG-PON): Transmission convergence (TC) layer specification, ITU-T
https://www.itu.int/rec/T-REC-G.987.3

[G.988] ONU management and control interface (OMCI) specification, ITU-T
https://www.itu.int/rec/T-REC-G.988

[G.9807.1] 10-Gigabit-capable symmetric passive optical network (XGS-PON), ITU-T
https://www.itu.int/rec/T-REC-G.9807.1

[G.Sup45] ITU-T (2009) GPON Power Conservation, Series G, Supplement 45, ITU-T
https://www.itu.int/rec/T-REC-G.Sup45

[RFC1492] An Access Control Protocol, Sometimes Called TACACS
https://datatracker.ietf.org/doc/rfc1492/

[RFC2236] Internet Group Management Protocol, Version 2, IETF
https://datatracker.ietf.org/doc/rfc2236/

[RFC2307] An Approach for Using LDAP as a Network Information Service, IETF
https://datatracker.ietf.org/doc/rfc2307/

[RFC2865] Remote Authentication Dial In User Service (RADIUS), IETF
https://datatracker.ietf.org/doc/rfc2865/

[RFC3046] DHCP Relay Agent Information Option, IETF
https://datatracker.ietf.org/doc/rfc3046/

[RFC3376] Internet Group Management Protocol, Version 3, IETF
https://datatracker.ietf.org/doc/rfc3376/

[RFC3810] Multicast Listener Discovery Version 2 (MLDv2) for IPv6, IETF
https://datatracker.ietf.org/doc/rfc3810/

[RFC4607] Source-Specific Multicast for IP, IETF
https://datatracker.ietf.org/doc/rfc4607/

[RFC6221] Lightweight DHCPv6 Relay Agent, IETF
https://datatracker.ietf.org/doc/rfc6221/

[RFC7761] Protocol Independent Multicast – Sparse Mode (PIM-SM): Protocol Specification, IETF
https://datatracker.ietf.org/doc/rfc7761/

[RFC8815] Deprecating Any-Source Multicast (ASM) for Interdomain Multicast, IETF
https://datatracker.ietf.org/doc/rfc8815/

[TR-101] Migration to Ethernet-Based Broadband Aggregation, BBF
https://www.broadband-forum.org/download/TR-101_Issue-2.pdf

[TR-156] Using GPON Access in the context of TR-101
https://www.broadband-forum.org/download/TR-156_Issue-2.pdf

[TR-383] Common YANG Modules for Access Networks, BBF
https://www.broadband-forum.org/download/TR-383_Amendment-2.pdf

[TR-385] ITU-T PON YANG Modules, BBF
https://www.broadband-forum.org/download/TR-385.pdf

[TR-384] Cloud Central Office Reference Architectural Framework, BBF
https://www.broadband-forum.org/technical/download/TR-384.pdf

[TR-402] Functional Model for PON Abstraction Interface, BBF
https://www.broadband-forum.org/download/TR-402.pdf

[TR-403] PON Abstraction Interface for Time-Critical Applications, BBF
https://www.broadband-forum.org/download/TR-403.pdf

[TR-451] vOMCI Interface Specification, BBF
https://www.broadband-forum.org/technical/download/TR-451.pdf

[TR-454] YANG Modules for Network Map & Equipment Inventory, BBF
https://www.broadband-forum.org/technical/download/TR-454.pdf

[VOLTHA] https://opennetworking.org/voltha/

[WT-413i2] SDN Management and Control Interfaces for CloudCO Network Functions, BBF

[WT-474] Subscriber Session Steering, BBF

[WT-477] Cloud CO Enablement– Access Hardware Disaggregation, BBF

[WT-489] Authentication of an Optical Network Unit (ONU), BBF

[X.733] Systems Management: Alarm reporting function, ITU-T
https://www.itu.int/rec/T-REC-X.733

Glossary

100GE	100 Gigabit Ethernet	IS-IS	Intermedia System - Intermediate System
10GE	10 Gigabit Ethernet	ITU-T	International Telecommunication Union (Telecommunications)
API	Application Programming Interface	LDP	Label Distribution Protocol
B/OSS	Business and Operational Support Systems (i.e. the whole IT stack)	LDP DoD	LDP Downstream on Demand
B2B	Business to Business	MAC	Media Access Control
B2B2X	Business to Business to … (i.e. wholesale)	MPLS	Multi-Protocol Label Switching
B2C	Business to Consumer	NBI	Northbound Interface
BBF	Broadband Forum	OLT	Optical Line Termination
BSS	Business Support Systems	OMCI	Optical Management and Control Interface
C-VLAN	Customer VLAN	ONF	Open Networking Foundation
CP	Control Plane	ONT	Optical Network Termination
CPE	Customer Premises Equipment	OSPF	Open Shortest Path First
CPU	Central Processing Unit	OSS	Operational Support Systems
CUPS	Control Plane / User Plane Separation	P2P	Point to Point
DHCP-RA	Dynamic Host Configuration Protocol Relay Agent	PIM	Protocol Independent Multicast
FTTA	Fiber To The Antenna	PON	Passive Optical Network
FTTB	Fiber To The Building / Basement	PPP-IA	Point to Point Protocol Intermediate Agent
FTTC	Fiber To The Curb	RIPv2	Routing Information Protocol version 2
FTTH	Fiber To The Home	S-VLAN	Service VLAN
FTTN	Fiber To The Node	SBI	Southbound Interface

FTTP	Fiber To The Premises	SDN	Software-Defined Network
FTTX	Fiber To The … (i.e. all fiber-to-the variants)	SW	Software
GPON	Gigabit PON	TIP	Telecom Infra Project
HSI	High Speed Internet	UP	User Plane
HW	Hardware	VLAN	Virtual Local Area Network
IGMP	Internet Group Management Protocol	VoD	Video on Demand
IPTV	IP Television	XGS-PON	10Gigabit Symmetric PON