Open Fixed Access Networks

Use Case Document

Date Jun 29, 2022

Document version: 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.com

Mauro Tilocca

Project Manager, Telecom Italia.

Mauro.tilocca@telecomitalia.com

Paolo Pellegrino

Fixed Access Innovation Project Manager, Telecom Italia.

paolo.pellegrino@telecomitalia.com

TIP Document License

Exhibit A

TIP CONFIDENTIAL

Change Tracking

Date	Revision	Author(s)	Comment
28/Jun/2022	v1.0	OFAN sub-group	Approved version for release

Table of Contents

1

Introduction

1. Introduction

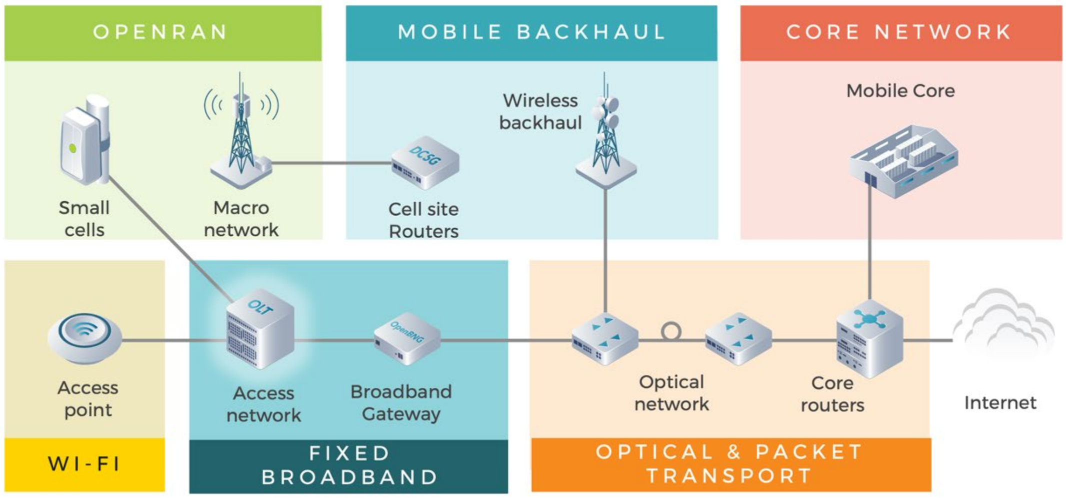

Figure 1 - TIP Fixed Broadband Project Group scope

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. As part of this work, a new sub-group has been formed, 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.

Against the backdrop of continued traffic growth driven by increasing consumer demand for immersive online experiences, fiber-based access providers are uniquely positioned to provide the fast, reliable, and cost-effective connectivity services their end users demand. In collaborating on this work, the sub-group members have identified the specific challenges that are most meaningful when building and scaling their access networks to meet these demands.

The evolution of the access network towards open technologies comes against the backdrop of similar initiatives across all areas of the access, transport and core networks. By engaging in this work, operators are hoping to replicate these successes and unlock benefits that will increase the level of innovation, efficiency and openness in their networks.

This Use Case Document (UCD) defines the overall scope of the project, and introduces the high-level characteristics of Open Access Terminal technologies. Since the ultimate objective for this work is to develop technologies that can be deployed into production networks, this work is expressed in the context of the specific use cases and deployment scenarios that are representative of how operators expect to deploy this tech in the real world.

2

Background and Objectives

2. Background and Objectives

The ultimate objective for this work is to drive the deployment of open, disaggregated and interoperable access technologies. In so doing, the sub-group participants are looking to address the issues described in the following sections.

2.1 Multi-vendor interoperability

Operators value the ability to deploy multi-vendor networks as a means to increase network diversity and drive innovation. However, while many open standards exist, operators remain unable to realise their vision – poor interoperability between supplier implementations means that operators instead find themselves locked in to homogeneous single-vendor networks, or obligated to follow a time-consuming process of tests and changes to achieve it.

As with other network segments, these interoperability challenges play out across multiple axes:

North-South interoperability	East-West interoperability
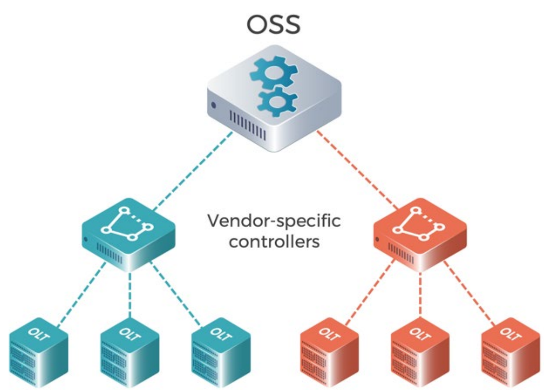	
The northbound programmable interfaces (APIs) implemented by the Optical Line Termination (OLTs) are based on closed interfaces and/or vendor-proprietary models.	The interfaces between OLT and Optical Network Termination (ONT) are based on open standards, but are often augmented with proprietary extensions that impair functionality in interoperability scenarios.

Figure 2 - Interoperability challenges with existing fiber access network technologies

While these challenges are not unique to the fixed access domain, the inherent characteristics of the access network mean this problem is amplified. In particular:

- There is a tight integration between the access network and the OSS, which is required to support subscriber fulfillment and assurance activities. This means that introducing new vendors and technologies requires costly, complex and time-consuming integration activities that extend deep into the IT stack.
- Swapping an OLT can require changing all ONTs deployed downstream, which is typically impractical to coordinate given ONTs are physically deployed in customer locations. This can mean operators are simply unable to consider new supplier solutions, or at the very least reduced functionality for affected customers during transition scenarios.

Put together, these issues mean that the operational costs to manage and operate the network are artificially high, and that operators face a very high cost of change. The Open Fixed Access Networks sub-group intends to address this by driving the use of truly open and interoperable interfaces for all network functions through clear and specific explanation of how each use case should be achieved using open standards.

2.2 Disaggregation

As with other parts of the network, ongoing consolidation of the supplier market has led to a high dependency on a decreasing number of suppliers. In turn, the lack of competitive tension between their suppliers means operators are facing increasing costs due to a lack of commercial competition, and have limited scope/opportunity for innovation due to a lack of technical competition.

The Open Fixed Access Networks (OFAN) sub-group believes that disaggregation has a vital role to play in creating a healthier and more competitive supplier ecosystem. Disaggregation in this context covers two key areas:

HW/SW disaggregation	Disaggregated architectures
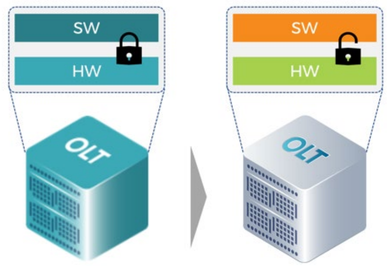	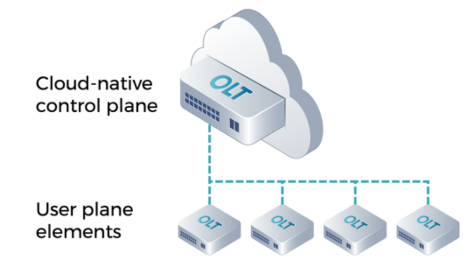
The separation of hardware (HW) based on 'whitebox' merchant solutions that can support a variety of different software (SW) solutions allows operators to mix and match different SW and HW products, and to change these choices of time	The use of modern, disaggregated architectures (e.g. [VOLTHA], [OB-BAA]) is a means to improve interoperability with existing and legacy networks, and enables operators to leverage existing and ongoing investments in cloud infrastructure

Figure 3 - Disaggregation in the context of open fiber access networks

HW and SW disaggregation is critical to allow operators to decouple the long-term (7-10 year) investments in hardware from the need to constantly evolve and innovate the SW capabilities those HW platforms support. As with other TIP initiatives, the sub-group expects that disaggregating their networks in this way will enable new suppliers to enter the market, significantly increasing choice.

Further architectural disaggregation is an exciting area of activity, with many recent developments emerging from the Open Networking Foundation (ONF), Broadband Forum (BBF) and other standards-defining organisations. The sub-group believes that adoption of these architectures also has a critical role to play in breaking the overall problem space into smaller pieces. In turn, this will lower the bar for new suppliers, enabling them to bring forth new products that are more innovative and cost-effective, and which can integrate and interoperate freely into existing (e.g. brownfield) as well as new build environments.

2.3 Operating efficiencies

Fixed network operators face a constant challenge to be more efficient in how they use their technical and human resources. In particular, the sub-group believes that there are opportunities in a number of areas:

Enable new operational paradigms

- Simplify fault detection by a newer and structured approach
- Adopt more effective and efficient network automation and SDN control paradigms
- Leverage Zero-Touch Provisioning (ZTP) to reduce the cost and complexity of commissioning and in-life maintenance activities
- Tie into business process orchestration to automate population of inventory systems, and to simplify customer fulfillment and assurance

Create more effective and efficient devices

- Reduce power consumption through a better utilization of the resources
- Increase the interoperability among OLT and ONTs.
- Consolidate multiple network functions into a single device (e.g. network/service functions, or use a single system to combine PON and other user-network interfaces (UNI) in a single device)
- Facilitate the technology evolution and migration through OLTs supporting multiple PON technologies.
- Unify and open the transceivers connected into the OLT PON ports to increase their compatibility

Embrace and adopt new technical architectures

- Reduce fault resolution times by implementing cloud native approach for virtualized software
- Adopt cloud-native technologies to enable convergence with other network functions (e.g. 5G, telco cloud, …)

2.4 Service opportunities

While the applications for fixed access technologies are many and varied, actual deployments often fail to realise the expected benefits:

- Fixed access products do not align well with industry drive towards distributed and converged network architectures, with 'lowest common denominator' capabilities in the access network providing only very limited functionality to support 'high touch' use cases (e.g. QoS, L3 functions)
- Many operators are yet to embrace fibre access technologies for the delivery of business services, or for other infrastructure needs (e.g. mobile backhaul). Indeed, the deployment and adoption of 10G based PON technology is slow, and commercials remain very poor given a lack of competition and services demanding such capacity

The Open Fiber Access Network (OFAN) sub-group intends to address these topics by clearly articulating how these use cases and deployment scenarios should be achieved using open and disaggregated technologies, providing clarity to the supplier ecosystem about which products will be most valued, and which solutions will be most widely deployed.

3

Solution Overview

3 Solution Overview

This section provides a high-level overview of the project and its positioning in the network. The intent here is to provide an illustration of the solution space, and to describe the key characteristics of the Open Fixed Access Networks (OFAN) solution(s). Further technical detail will be provided in future Technical Requirements Documents.

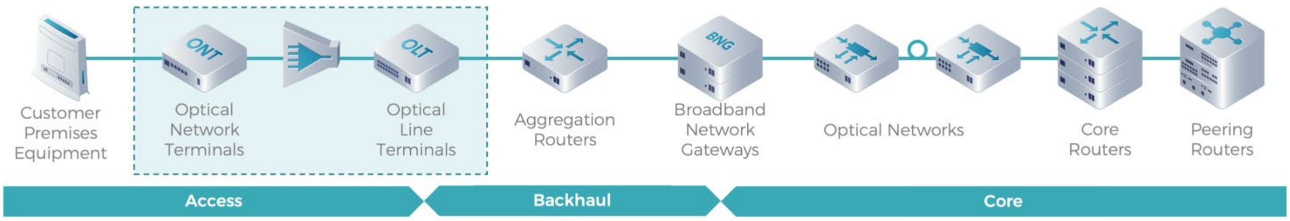

Figure 4 - Open Fiber Access Networks in the context of the end to end network

3.1 Defining the scope of our work

The primary scope of the project is the Optical Line Termination (OLT) according to ITU-T Recommendations and the management of the attached Optical Network Termination (ONT) by means of the ONU Management and Control Interface (OMCI), as well as interfaces towards a northbound Software Defined Network (SDN) controller.

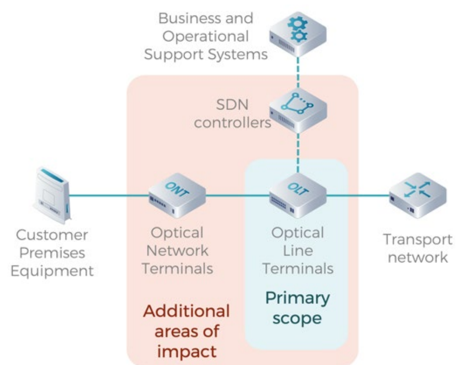

Figure 5 - Open Fiber Access Networks scope

While the primary focus of this initiative is to develop OLT technologies, we do acknowledge that meeting our interoperability objectives will necessarily impact many other areas, including integration points with adjacent ONT/ONUs, as well as integrations and use cases/processes implemented by northbound controllers. These topics are explored in more detail in the following sections.

3.1.1 Impact on the ONT Ecosystem

We believe that the availability of OLTs that implement open interfaces in a consistent manner will open up the ecosystem to give operators more choice about what ONTs they want to deploy. However, we also recognise that there is a corresponding need for a diverse supply of ONTs that can be deployed into such an interoperable network. As such, while we don't intend to develop detailed requirements that explicitly relate to ONTs, we do expect to cover the following areas:

- Clearly defining the OMCI interface requirements for the OLT, which in turn guide for the corresponding ONT interfaces.
- Defining behaviours based on standard OMCI Management Entities to mitigate the need and scope for proprietary extensions.

We also believe that testing of open and disaggregated OLT products with commercially available ONTs has an important role to play, and we look forward to testing the products and solutions vendors will develop, with a particular focus on validating interoperability between a wide range of OLT and ONTs.

3.1.2 Impact on the SDN Controller Ecosystem

Control and management of OLTs is a critical topic, covering two main areas:

- Integration with a domain-specific SDN controller for the purposes of element management (e.g. fault/performance management, alarms, diagnostics).
- Integration with the high levels of the IT stack (BSS and/or OSS) for the purposes of subscriber fulfilment, assurance and provisioning.
- We expect that our scope should cover the first point, considering that we will clearly define the standardised interfaces (referencing appropriate protocols and models) that should be used to control and manage the OLT. In turn, these requirements will also imply corresponding SBI requirements that an SDN controller should implement to manage an open OLT.

Regarding more complete alignment of the interfaces towards the B/OSS layer, we believe that more extensive standardisation work is needed in this area. As such, we do not expect to describe the SDN controller NBIs. However, we believe this work is important to truly break down the barriers to multi-vendor access networks, and welcome opportunities to collaborate further with other organisations to cover this topic in the future.

3.2 Characteristics of an open and disaggregated OLT

The OLT is the key element of this project. It consists of a number of functional components, which are provided by a combination of whitebox HW coupled with Network Operating System (NOS) SW.

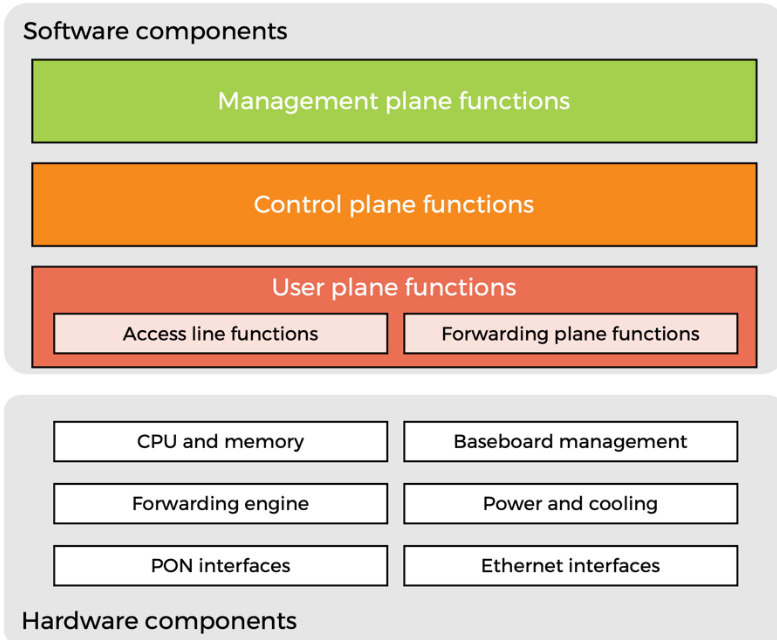

Figure 6 - Characteristics of an open and disaggregated OLT

These components are classified according to their functional role as follows and in accordance with the Broadband Forum reference architecture [TR-384]:

- Whitebox hardware, which provides physical interfaces for traffic and management, a forwarding element to process traffic, CPU and memory to host SW functions, and supporting power/environmental functions
- User plane functions that handle the forwarding of traffic between interfaces
- Control plane functions that interact with adjacent network elements to control user plane behaviours
- Management plane functions that interact with controller and operators to program device behaviours

The following sections describe key characteristics of the solution, and are intended to provide only a very high-level overview for contextual purposes. Subsequent technical requirements documents will provide further, authoritative, details of each component.

3.2.1 Hardware characteristics

The physical OLT HW provides network interfaces as well as compute and memory resources that support the SW running on top of it. Since the objective for this project is to deliver simple and interoperable solutions, we expect HW solutions to adopt ‘pizza box’ form factors with standards-based network interfaces in preference to chassis-based systems with proprietary connectors/interconnects. In some applications the HW solution can adopt also a SFP form factor in order to transform a port of a general purpose switch in a single port OLT device.

If more pizza box OLTs are needed in a central office, they have to be aggregated by redundant Top-of-the-Rack switches in order to have a single high capacity link towards the transport network.

OLTs are frequently deployed into hostile physical environments, often with constraints on power availability and limited cooling capacity. As such, we expect HW solutions to be optimized for low power consumption and heat dissipation, and with minimal need for ongoing maintenance activities.

The interface specification will vary according to each deployment scenario, and will be explored further in the technical requirements phase. In general terms, devices will need to support both GPON and XGS-PON interfaces towards the access network, and multiple 10G/100G interfaces towards the transport network. A limited number of Ethernet ports to service downstream access connections would be valued.

As operators are currently moving towards XGS-PON, but GPON still remains for years, in order to optimize space occupancy and power consumption, it is important to also look at combo solutions, where GPON and XGS-PON are on the same port.

3.2.1.1 PON technologies

In terms of PON MAC technology, ITU-T based standards will be followed. At the time of writing, GPON technology is well established in most regions and XGS-PON is the next massive deployment which is already in planning or early deployment stage by most network providers. As such, we expect that operators deploying open OLTs will need to support both GPON and XGS-PON connections for the visible future. Three main approaches exist:

- Single technology OLTs which only support GPON or XGS-PON technologies at all ports (but not both).
- COMBO-PON OLTs, supporting both GPON+XGS-PON at all ports (but do not support single technology at any port).
- Flexible PON technologies, supporting GPON, XGS-PON or Multi-PON modules (MPM) in a flexible way at a PON port basis.

We believe the most interesting and cost-effective approach is to start deployment with Flexible PON OLTs using GPON pluggables at day 1 and replace them with MPMs at a PON port basis as XGS-PON technology demand grows.

3.2.2 Software characteristics

As described earlier in this section, the solution consists of a number of SW functions, which can be further broken down as follows:

Management Plane Functions		• Management and control of downstream ONTs • Northbound interfaces towards orchestrators and other elements that control the OLT • Alarms, diagnostics and logging interfaces • Performance measurement and monitoring
Control Plane Functions		• Broadband Forum Helper functions (e.g. PPP-IA/DHCP RA) • Integration with transport network (e.g. routing protocols)
User Plane Functions	Forwarding plane	• Quality of service towards PON (access-facing) and Ethernet (core-facing) domains • Encapsulation and forwarding of customer traffic
	Access line	• PON interfaces towards end customers • Dynamic Bandwidth Allocation • Quality of Service on shared access lines

There are two ways these functions can be provided:

- Physical Network Functions (PNF) that are deployed locally on the HW element, using local CPU and memory resources provided by the whitebox HW
- Virtualised and provided by an external server, or deployed in a cloud environment

While most control and management plane functions can be deployed as local or remote/virtual network functions, many user plane functions must inherently be deployed locally on the whitebox given the level of integration required between the SW and forwarding element. Accordingly, we can consider that control plane functions may be deployed locally on the PNF, or more centrally in a virtual/cloud environment:

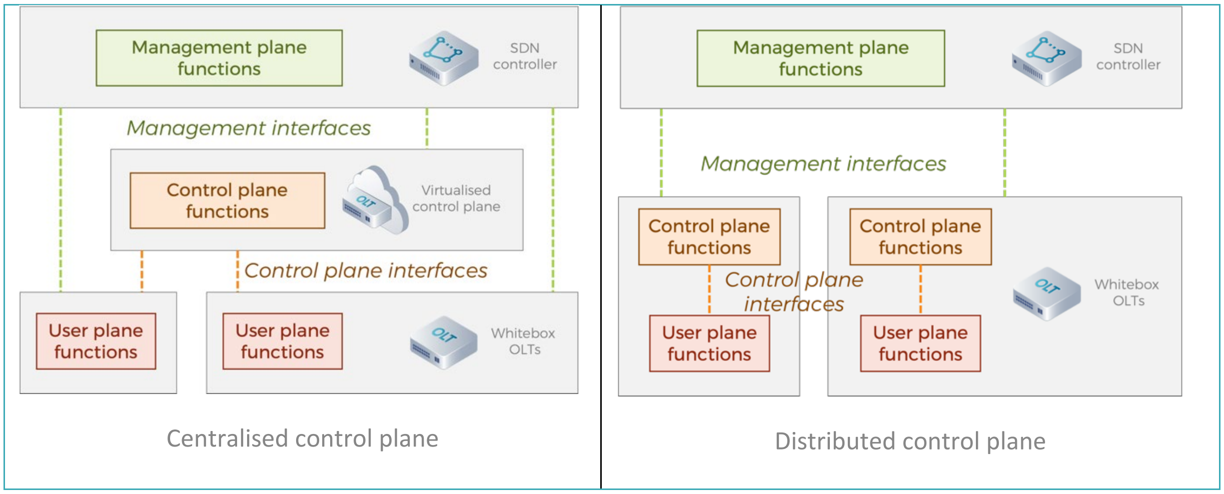

Figure 7 - Centralised vs distributed control plane models

The schematic above depicts the functional split between PNF and virtual network functions. A centralised deployment of the control plane in a cloud environment is expected to be more efficient in high density FTTH deployments, whereas a distributed deployment within the internal resources of the whitebox OLTs is expected to be more efficient in low density FTTH deployments or in those scenarios where data centre resources are not available.

4

Deployment Models

4 Deployment models

While the architecture and design of each operator's network is ultimately unique, these deployments typically share many common characteristics. For the purposes of this document, we can consider the following aspects:

- Deployment model, focusing on where the OLT is physically located, how it integrates with the customer and core network environments, and how it scales
- Services model, considering what end user services the OLT is supporting, what functions it provides in support of those services, and how those functions are provided

For the purposes of this section, we explore each area individually. We provide a description of the different approaches most commonly adopted by operators, along with an assessment of the benefits and considerations for each approach.

Since these aspects are ultimately inter-related, later sections will group them into higher order use cases. These use cases represent real world deployments that operators expect to make with fixed access technology, and will form the basis for the onward definition of technical requirements. The same use cases will also be used to structure onward test and validation activities.

4.1 Physical Deployment Models

Here we consider the different physical architectures in which we can deploy OLTs. We focus on four areas:

- A macro view of the network, describing the different types of FTTX network that can be deployed
- A description of the different locations where OLTs can be deployed
- A site-level view of the network, describing how OLTs can be deployed in each location
- A device-level view of the network, describing the characteristics of the different OLTs that would be deployed according to the above models

4.1.1 FTTX Network Types

There are multiple different types of FTTX network currently being deployed. The key difference between these network types is where the customer-side demarcation is, depicted as follows:

Dedicated access variants

Fiber-T-The-Home (FTTH)	Fiber-To-The-Premises (FTTP)

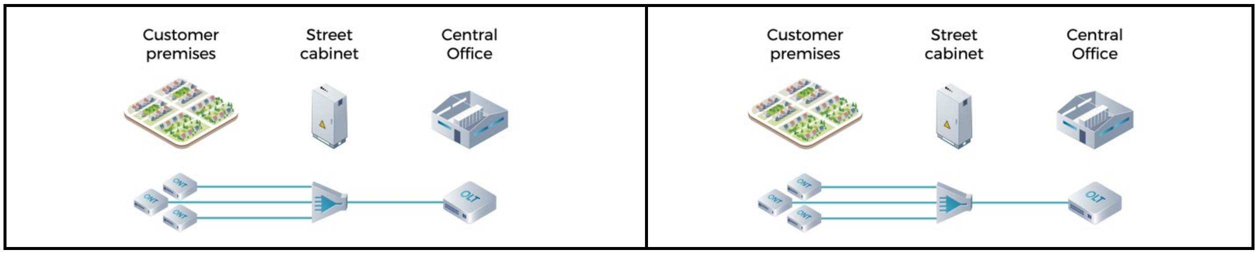

Figure 8 - Dedicated access deployment variants

Shared access variants

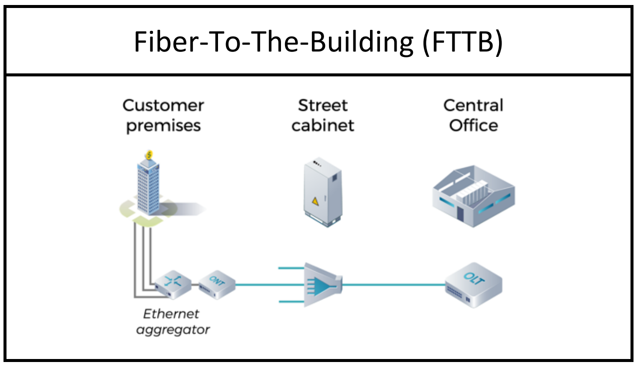

Figure 9 - Shared access deployment variants

Infrastructure variants

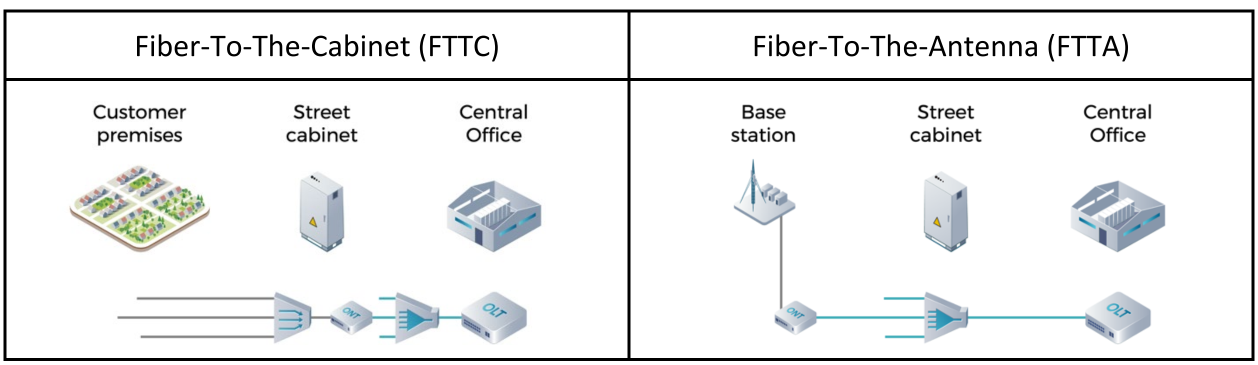

Figure 10 - Infrastructure deployment variants

In all cases, the scope of this project extends only as far as the optical termination device. In the case of end-customer services (FTTB, FTTH, FTTP), this demarcation of scope occurs at the ONU where the broadband interface is presented as a single electrical Ethernet port. In cases where single-box CPE is supplied i.e. a Home or Business Gateway with an integrated ONU with Ethernet switch and Wi-Fi, then we are only concerned with the functionality of a single (primary) Ethernet port.

While the specifics of each service offering are beyond the scope of this project, we do expect open OLTs to be used to deliver services in scenarios. Therefore, the sections that follow provide a brief overview of each scenario for introduction and context.

4.1.1.1 Dedicated access - Fibre-to-the-home / premises (FTTH / FTTP)

These models describe where the PON is demarcated directly in the end customer environment, and where the ONU integrates directly to the CPE. These are two models here:

- Connectivity services provided to residential users. For the purposes of this document we term this model as FTTH (fibre-to-the-home)
- Connectivity services provided to business users. For the purposes of this document we term this model as FTTP (fibre-to-the-premises)

While the home (FTTH) and business (FTTP) variants are structurally very similar, the service end customer service offering in each case is typically quite different:

	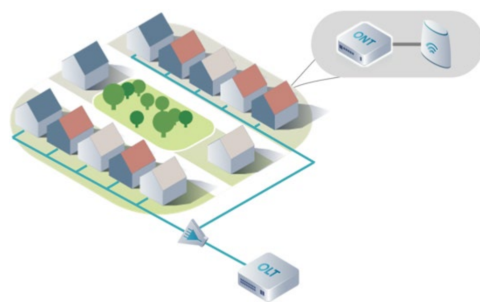 FTTH (residential)	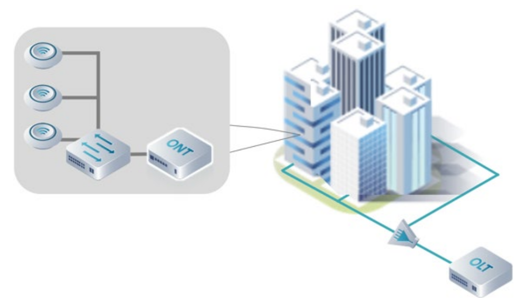 FTTP (business)
CPE	CPE is likely to be an ONU with a Home-Gateway. In some cases the ONU and CPE may be integrated in a single unit.	CPE is likely to be a dedicated ONU, coupled with a Business Hub or Ethernet switch, providing connectivity for a more extensive in-building network
SLA obligations	Fault repairs/fix are typically best endeavours, with limited scope for financial compensation	Service restoration is typically guaranteed within an agreed time period

These differences mean that, while the underlying technologies used to deliver FTTH and FTTP may be very similar, the associated service configurations can be very different. These issues are explored further in section 4.3 later in this document.

4.1.1.2 Shared access - Fibre-To-The-Building (FTTB)

In the fibre-to-the-building (sometimes known as 'fibre-to-the-basement'), fibre is terminated within a building by an ONU, and then a broadband service is distributed to numerous subscribers around that building via another access technology (typically copper-based P2P Ethernet). Subscriber fulfilment, assurance and billing are provided by the in-building access platform, rather than by the FTTB platform.

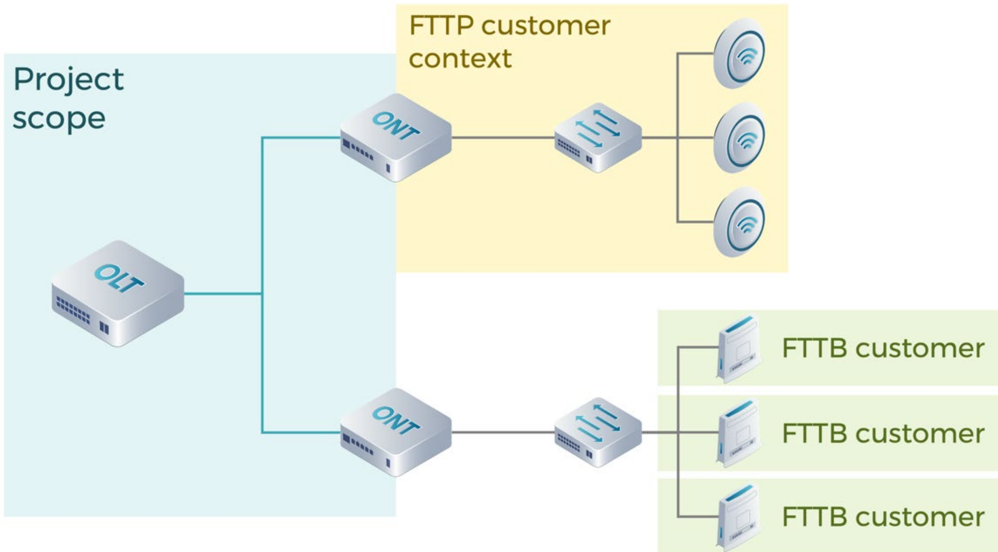

Figure 11 - FTTP vs FTTB

The FTTB topology is superficially similar to the FTTP topology - they both use a single PON connection to service multiple units in a single building, with an Ethernet switch providing connectivity for many users within that building. However, the service architecture is quite different:

- In the FTTP model each PON service represents a single account for billing and support purposes, which provides access for all users in the building
- In the FTTB model the PON service is used to provide the CSP infrastructure, and instead the in-building access technology provides a unique account for each unit/dwelling

4.1.1.3 Infrastructure access - Fibre-to-the-Node / Fibre-to-the-Antenna

In the Fibre-to-the-Node and Fibre-to-the -Antenna cases fibre connections provide essential backhauling transmission capacity. Possible examples are mobile radio access infrastructures in which the fibre is typically terminated into an optical node or a gateway device located in proximity of an antenna site.

4.2 OLT Deployment Locations

The previous section described the different types of FTTX network, which are differentiated by the location of the customer demarcation and the end customer proposition it supports. The purpose of this section is to describe the different locations and environments where an OLT can be deployed. The sections below will then describe the specific considerations for the devices that would be deployed into these locations.

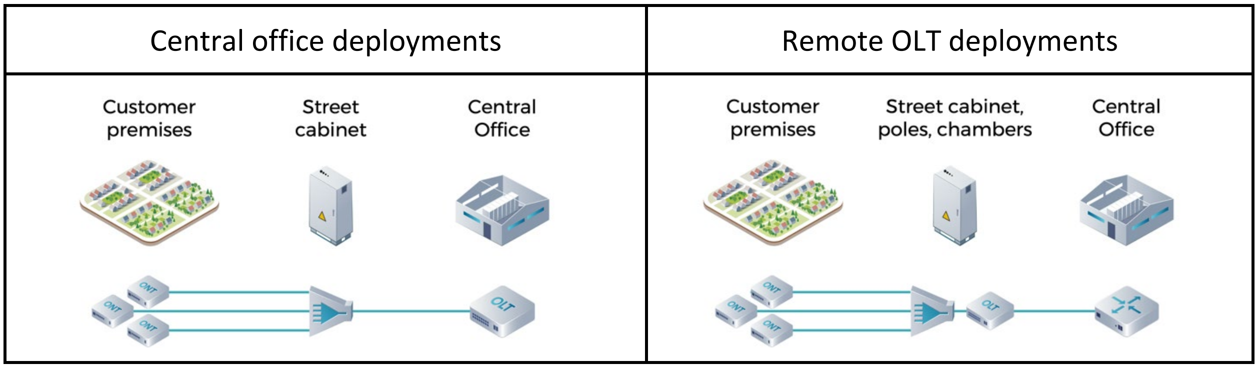

Figure 12 - Indoor and outdoor OLT deployment locations

Note that we consider the location of the OLT to be largely agnostic to the network types described in the previous section. That is to say, any of the FTTX variants described above can be delivered using OLTs deployed to Central offices, or OLTs deployed in remote locations.

4.2.1 Central office deployments

This model describes the case where an OLT is deployed to a physical building operated by the CSP. Central office environments vary considerably according to local operating conditions, which can be divided into two general use case categories:

- Data-centres - where the OLT is mounted into a rack and where forced-air cooling allows power/port densities to be achieved that cannot be achieved using passive means. Accordingly, data-centres can support a variety of different types of OLT
- Exchange-based – where the OLT is mounted into a rack and where no (or limited) forced-air cooling is possible. Exchange-based OLTs are commonly mounted in 'back-to-back' configurations that mean the depth of each device must be limited to 300mm. In this case, it is expected that a 'hardened' pizza box type OLT that has been optimised for low power consumption and heat dissipation would be needed.

4.2.2 Remote OLT deployments

This model describes the case where the OLT is deployed in a remote environment, typically housed within some sort of street furniture. Common considerations for all such deployments include the need for environmental hardening (e.g. heat tolerance, IP rating) and the need for space-optimised physical designs. Remote OLTs can also be divided into two categories:

- Remote Cabinet OLT – where the OLT is mounted inside a street cabinet which may have either zero or some cooling capability (say active heat exchange doors). It is expected that the OLT type installed at this location would be similar to that installed in an exchange with limited or no aircon available.

- Remote Underground OLT – where the OLT is located within a hermetically sealed container.

4.2 OLT Deployment Topologies

For any given deployment location, we consider that OLTs could be deployed into one or two topologies as depicted below:

Standalone OLT	Clustered OLT
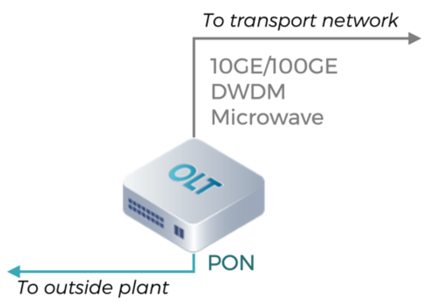	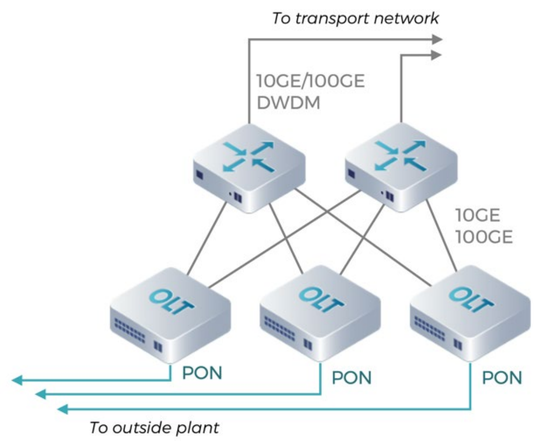

Figure 13 - Standalone vs clustered OLT deployments

The standalone OLT topology sees the OLT directly connected to a backhaul link that connects it to the upstream transport network. This topology is most efficient for low density sites where minimal user/traffic growth is expected, and where capacity demands can be met by a single OLT. However, where needed it is possible to deploy multiple OLTs to a single site, considering that each OLT will operate as a unique entity in such a scenario.

Meanwhile, the clustered OLT topology sees the OLT deployed alongside a dedicated aggregation infrastructure, and the aggregation infrastructure terminates backhaul connections towards the transport network. This topology can accommodate a large number of OLTs while also making efficient use of backhaul connectivity. It can also allow for some degree of functional/service separation between OLT and aggregation nodes. However, it also requires the deployment of additional equipment, which can be problematic for constrained deployment locations.

4.3 OLT Device Types

The following sections describe the key physical characteristics of open OLT HW devices.

4.2.1 Form factor

There are two main form factors commonly used to build OLT devices today - a chassis-based OLT and a pizza box approach. The figure below outlines the key characteristics of each

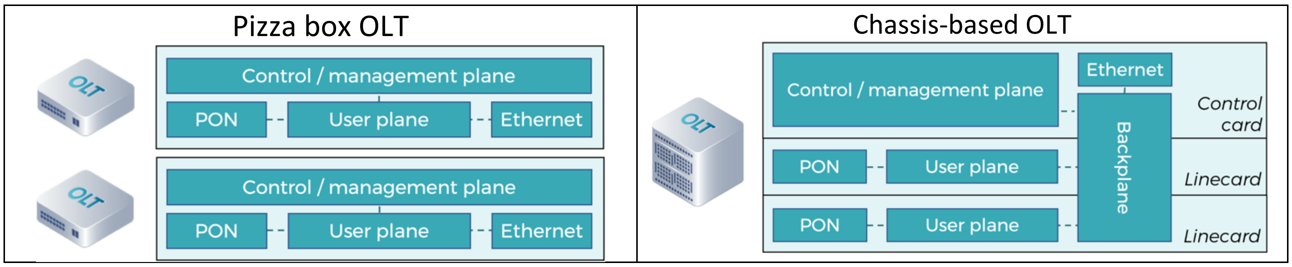

Figure 14 - Pizza box vs chassis OLT form factors

The main difference is the switch engine and fabric interfaces that allows the line cards to be plugged into the backplane. Obviously for a chassis-based system that is sparsely populated with line cards then maybe a pizza box based OLT design would be more efficient. However, chassis-based designs are more efficient from a power consumed per port basis when they are more highly populated with line cards.

Note that while chassis-based systems can deliver significant scale and port density, there are few open standards that define the interface between chassis and linecard components, and typically this interface is proprietary for each manufacturer. Consequently, we expect that a self-contained 'pizza box' form factor is the most desirable and interoperable configuration for open OLT HW products.

An alternative to pizza box and chassis-based OLT is the SFP form factor. It can allow an additional degree of flexibility in designing the PON network architecture for sparse FTTH deployments or new applications. In this case a fabric or even a small and simple switch device in the network can be very easily transformed, for a certain number of its ports, in an OLT. We believe this could be an interesting option for future use cases to be explored during subsequent phases of the project.

4.2.2 Industrial design

Depending on the OLT deployment location, the HW may be subjected to harsh physical operating conditions. Accordingly, we can consider three types of OLT HW will be required:

- A datacenter-optimised configuration, which can be optimised for density
- A CO-optimised configuration, also optimised for density, but which is also temperature hardened or thermally passive (accepting that there may be some performance/capacity trade-offs)

- Remote-optimised configurations, which service lower densities and are instead optimised for environmental hardening (e.g. thermal performance, IP rating) and can be installed in street cabinets or even underground.

4.3 Service models

Operators use fibre access technologies to provide a wide range of connectivity services to a wide range of end users. This will typically include:

Service types	End customer types
Internet ('HSI') connectivity Multi-play (e.g. voice/video/data) services Layer 3 connectivity into VPNs Layer 2 connectivity into VPNs Backhaul for legacy/3GPP access Multicast VLANs/PIM routing	Residential (B2C) SME (B2B) Enterprise (B2B) Wholesale (B2B2X) Internal/infrastructure users

To deliver these services, the network needs to provide the following service functions:

- User identification, authentication, accounting and authorization
  - In the case of FTTX, this covers both the CPE and the ONU
  - This also includes things like IP address allocation
- Layer 2 transport
  - This includes transport of S-VLAN and C-VLAN tagged traffic
  - Quality of service and traffic differentiation
  - This covers a full spectrum of approaches to traffic management
- Upstream and downstream scheduling onto the PON
  - 802.1P within the transport and home networks
  - DSCP/EXP differentiation within the transport and subscriber management domains
- IP routing
  - This covers IPv4 and IPv6, unicast and multicast
  - LDP DoD protocol for MPLS
  - Routing protocols (IS-IS, OSFP…) in order to exchange MPLS labels
  - Static routes
  - IPoE services in VLANs (VRF) where OLTs are gateways of an IP subnetwork for specific services (ex: VoD)

  - RiPv2 in VRFs for dynamic routes support in residential gateways
  - IGMP proxy/snooping for IPTV services in multicast VLANs
  - PIM routing for IPTV

While these functions are common to all service types and use cases, certain use cases may require specific functions (e.g. precision timing/synchronization). These specific functions are described in the relevant use cases outlined below.

Broadly speaking there are two architectures operators can adopt to deliver these services:

- Layer 2 deployments, where the OLT provides only aggregation and VLAN switching functions, and higher order service functions (e.g. IP routing, MPLS transport) are provided by the upstream transport network
- Layer 3 deployments, where the OLT provides both aggregation and service termination functions

Each architecture provides the service functions in a different way, as depicted in the sections below.

Each architecture can be deployed using a distributed OLT (CP and UP locally on the same HW), or in a cloud environment involving a centralised CP (with CP functions separated from the UP PNF). These models are introduced in section 3.2.1 above, and depicted in Figure 7.

The sections that follow will describe each of these architectures, including:

- How each of the service types described above are delivered according to that architecture
- Variations / options that exist within that architecture
- Any relevant benefits and considerations (pros and cons) of the architecture

4.3.1 Layer 2 deployments

In this deployment model the OLT performs layer 2 functionalities applied on the traffic flows identified at the ONT U interface and at the OLT V interface. Basically, the GPON/XGSPON access system behaves as an access aggregation device capable of upstream and downstream ethernet traffic classification, scheduling and forwarding, MAC learning and bridging, VLAN tag manipulation. Furthermore the system is provided with multicast handling capability to enable an efficient distribution of video services to the final users (e.g. IPTV).

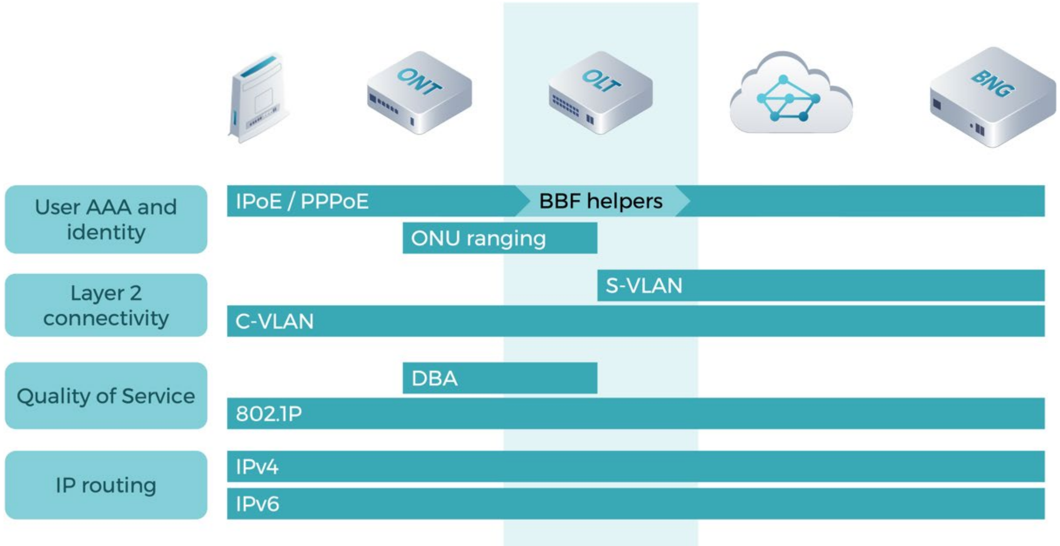

Figure 15 - Layer 2 service model

Typical scenarios employed by operators are the N:1 and 1:1 VLAN topologies as described in BBF TR-101 and TR-156.

4.3.2 Layer 3 deployments

In this scenario, OLTs provide, besides PON access, Layer 3 and MPLS functionalities in order to process the aggregated traffic from the FTTH-PON network to the BNGs and the transport network. Each FTTH service may be transported using routed VLANs (IP over Ethernet) or MPLS L2VPN tunnels, thus LDP, static routes and routing protocols (IS-IS, OSPF, RiPv2…), as well as IGMP proxy/PIM routing for IPTV service are required. Layer 3 deployments can coexist with a layer 2 deployment in the same OLT, in the sense that a service can follow a layer 2 approach, such as a High Speed Internet service, and other services can use a layer 3 architecture, like IPTV and/or VoIP.

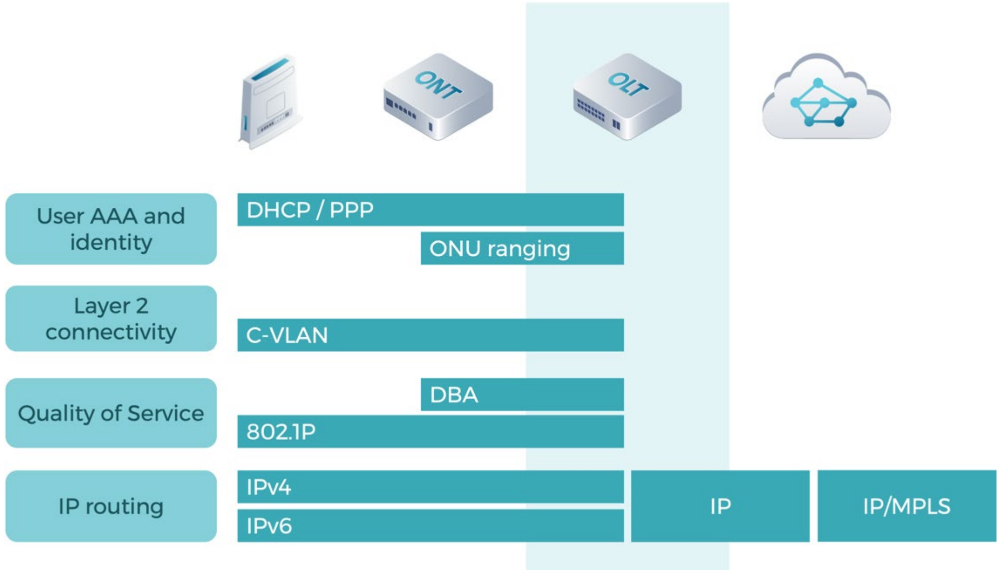

Figure 16 - Layer 3 service model

These functionalities may be implemented inside the OLT whiteboxes or in a spine pre-aggregation stage using devices such as DCSG, before the traffic reaches the BNGs or the aggregation network.

4.3.3 Cloud deployment

Layer 2 and Layer 3 service models can also be deployed in case of cloud deployment. As reported in Figure 17, control plane functions can be in this case terminated into the virtualized control plane (see Figure 7). In the examples reported in the figure all the functions are terminated into the virtualized control plane, but operators can decide to terminate some of them in the virtualized control plane, while others can remain in the physical OLT.

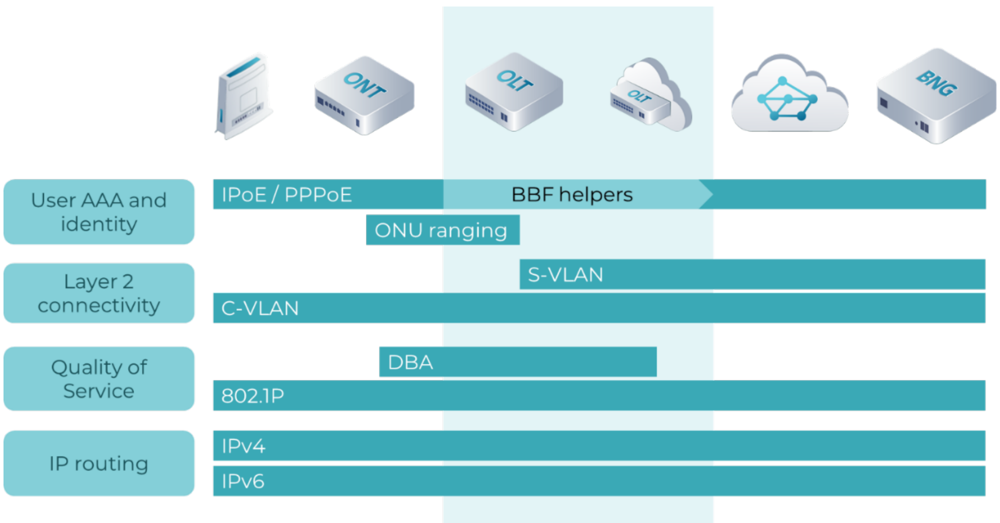

Figure 17 - L2 service model implemented using a centralised / cloudified control plane

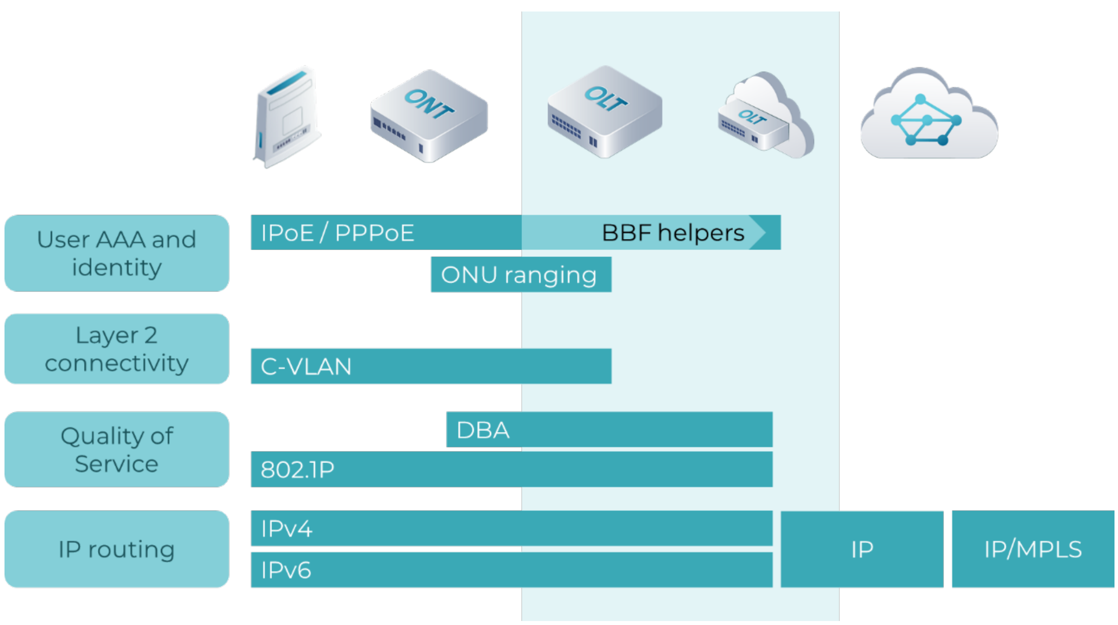

Figure 18 - L3 service model implemented using a centralised / cloudified control plane

5

Use Case Definitions

5 Use Case Definitions

This section draws together the unique combinations of the solution characteristics and deployment models introduced earlier in the document to form use cases. Each use case represents a common scenario where operators expect to deploy an open fiber access network device.

Use Case ID	Name	Description
OFAN-UC1	Exchange-optimised OLT	A disaggregated OLT with pizza box form factor, which is optimised for deployment in an Exchange (as described in section 4.2.1 above)

At the time of writing, the sub-group has defined only a single use case, but expects to add further use cases over time based on contributions from new operators, as well as the ever-evolving industry landscape. For completeness, these use cases will draw upon the same definitions and context outlined earlier in this document.

5.1 Use Case: Exchange-Optimised OLT

The Exchange-Optimised OLT use case describes an OLT with a fixed ('pizza box') form factor. For this use case the OLT is typically deployed into a Central Office environment, but may also be deployed into remote locations such as active street cabinets. Variants offering industrial grade temperature operation may be required to accommodate such remote deployments.

Each OLT provides between 16-64 ports of Flexible PON, and 10G/100G Ethernet ports for uplink, along with ports for local and console management. Multiple OLTs are frequently deployed to a single location for scalability purposes, and will typically be deployed in a clustered configuration. Multiple HW SKU variants are expected; each variant will address a unique scalability target, with systems providing dense PON aggregation and 100G uplink intended for Central Office deployments, and less dense/10G-optimised HW designs more suited for deployment in remote street cabinet locations. Further guidance regarding variants will be provided in the subsequent Technical Requirements Document.

Central Office OLTs optimized primarily for the delivery of FTTH and FTTP services, but may also be used to support small numbers of additional services. In particular, on-board Ethernet ports may also be used for other infrastructure uses (e.g. subtending or backhauling of other services). The Central Office OLT provides by default the layer 2 service model, and may optionally be configured for layer 3 services for some deployment scenarios.

6 References

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

[OB-BAA] https://www.broadband-forum.org/open-broadband/open-broadband-software/open-broadband-broadband-access-abstraction-ob-baa

[TR-384] Cloud Central Office Reference Architectural Framework, BBF

Glossary

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

IGMP	Internet Group Management Protocol	**VLAN**	Virtual Local Area Network
IPTV	IP Television	**VoD**	Video on Demand
IS-IS	Intermedia System - Intermediate System	**XGS-PON**	10Gigabit Symmetric PON

TELECOM INFRA PROJECT