

Vodafone L2 WAP Requirements



Vodafone Fixed Access Centre of Excellence

15 July 2015

C1 – Public



Document Information

Contacts: Gavin Young Gavin.Young2@vodafone.com
 Alexandre Serot Alexandre.Serot@vodafone.com

1. Contents

1. Contents	1
2. Executive Summary.....	2
3. Introduction	5
4. Architecture & Technology.....	6
4.1 Terminology	6
4.2 Architectural Standards Context & Reference Models.....	7
4.3 Ethernet VLAN Connectivity Options for L2 WAP	10
5. Network Characteristics	18
5.1 Network Interconnect	18
6.2 CPE/Modem Choice	21
6.3 Bandwidth, Traffic Prioritisation & Quality of Service	27
6.4 Multicast	31
6.5 Security & End-User Identification	33
6. Service-Wrap Characteristics	35
6.1 B2B & Portal Interfaces	35
6.2 Provisioning	36
6.3 Assure.....	37
6.4 SLA/SLG	39
7. Evolution Considerations.....	42
7.1 Access Transmission System Evolution	42
7.2 The Impact of Virtualisation.....	42
8. Prospects for L2 WAP Product Alignment Across Markets	44
9. Abbreviations	46
10. References.....	48



2. Executive Summary

This document describes Vodafone's requirements for a Layer 2 (Ethernet) Wholesale Access Product (L2 WAP) that is fit for purpose for offering competitive and differentiated retail services. It summarises the key L2 WAP functionalities necessary. It brings together requirements, material from Broadband Forum, MEF and NICC standards together with Vodafone's experience of using L2 wholesale products in a number of different local markets in order to articulate L2 WAP "best practise". It should be read in conjunction with Vodafone's submission to BEREC's consultation titled Vodafone's response to BEREC's Consultation on Common Characteristics of Layer 2 Wholesale Access Products in the European Union, 15 July 2015.

The architectural models defined in Broadband Forum TR-145/TR-178 documents together with the technical capabilities covered in the other referenced Broadband Forum Technical Reports and NICC standards provide L2 WAP network infrastructure providers with a complete set of solutions for meeting all the key requirements of Alternative Network Operators (ANOs) over a common, scalable network architecture.

There are a number of key capabilities of L2 WAP which have a profound impact on its utility to ANOs such as Vodafone. The product attributes in these areas influence and bound the ability of ANOs to develop competitive and differentiated products for end-user customers. A L2 WAP regulated product should enable operators to control and deliver the best possible customer and offer innovative new services.

The key product characteristics and associated best practise highlights are in the following areas:

- Network Architecture & VLAN configuration
 - Support for S-Tagged UNI with a VLAN per service at the residential customer UNI
 - Support for S-Tagged or Port Based UNI for business customer UNI.
 - Support for S-Tagged or S+C Tagged on the same NNI.
 - Minimum of 4 VLANs per residential customer with support for the same set of VLAN IDs on all UNI
 - Full flexibility for the access provider to define the ranges and policy for VLAN allocation on the UNI and NNI.
 - Access to an N:1 VLAN for multicasting
 - Unicast services can be identified uniquely using only VLAN IDs (not MAC address)
 - VLAN per customer at NNI and UNI for business customers
- Network interconnect location options
 - Ethernet Interconnect at CO Level
 - Parent-Child CO aggregation hand-over option and the possibility to request lower level hand-over at a street cabinet
 - 1G, 10G and resilient handover option (inc. n*1G & n*10G)
- User equipment options – CPE/modem choice
 - Wires-only option for ANO to provide own branded & integrated NTE/modem/ONT + router CPE
 - CPE interoperability requirements follow BBF standards and certification approach



- CPE interop test/validation environment to facilitate expanded “white list” of approved CPE
- Bandwidth, traffic prioritisations and QoS
 - Bandwidth profile available up to the maximum speed achievable by the DSL physical layer transmission system (including rate-adaptive systems)
 - Ability to request new bandwidth profiles for FTTH transmission systems
 - Choice of bandwidth profiles at least equal to those used by the L2 WAP Network Provider for their own retail services with the ability to request additional bandwidth profiles
 - Minimum of 4 levels of traffic prioritisation based on p-bits
 - Minimum quantified guaranteed throughput rates for upstream and downstream traffic (CIR from UNI to NNI) for over-booked and uncontended services
 - QoS SLA targets for each VLAN from UNI to NNI, defined by Frame Loss Rate, Frame Delay and Frame Delay Variation measures
- Multicast
 - Multicast Frame replication functionality for Local and Regional handovers on dedicated N:1 VLAN
 - IGMPv3 snooping for end-user access control of multicast in accordance with BBF TR-101 (& MLDv2 for IPv6)
- Security & End-User Identification
 - Customer Identification by Access Node and physical port identifiers (which can then be communicated via DHCP option 82 or PPPoE IA) - VLAN identifiers are an alternative only in situations where they provide unique customer identification co-ordinates
 - MAC address anti-spoofing - duplicate MAC address detection and rejection of traffic from duplicate MAC address sources
 - Control and policing of IGMP rate for N:1 VLANs
 - Rate limit Layer 2 broadcast
- B2B & portal interfaces
 - Automation of all key process interactions between L2 WAP Network Provider and ANO via B2B and portal interface options
 - SLA on systems availability and response time
- Provisioning process
 - Clear processes for provision and handling of errors or changes
 - Individual customer order progress reported regularly during all phases of provisioning
 - All key customer provide status milestones automatically notified via B2B interface and portal
 - Interconnect planned for growth, minimal upgrade impact on live traffic
 - Planning information on NGA rollout/coverage (by L2 WAP Network Provider) provided on a timely and regular basis
 - Ability to select and configure DSL line profiles/parameters and DLM stability thresholds
 - Option to self-provide the fibre drop within the multi-tenancy building for FTTH



- Assure process
 - Automated capability for confirming actual configuration of provisioned parameters (line profile, VLAN configuration etc.) and performance measures
 - Diagnostic capability for L1 and L2 testing (including loopbacks) for analysis of end user connections and for interconnects
 - Access and management to end user CPE from ANO network via industry standards based in-band methods
 - All key customer repair status milestones automatically notified via B2B interface and portal
- SLA/SLG
 - SLAs & SLGs should be defined for each of the main elements of the life cycle of services including at least, Ordering, Provisioning, Service Availability & Fault restoration
 - SLAs & SLGs should be defined for the electronic platform used to interface with the L2 WAP Network Provider, including availability and response times.
 - SLAs should be aligned with end-user requirements
 - SLGs should be high enough to incentivise compliance with the SLAs by the L2 WAP Network Provider, preferably with no penalty caps and with a right to claim for additional losses above the level of SLGs
 - SLAs & SLGs should apply per fault/event/line/circuit - not in aggregate for average performance
 - Payment of penalties should be pro-active /automatic – The ANO shouldn't have to measure it or ask for it
 - The L2 WAP Network Provider should provide reports on actual performance against SLAs – ANOs should have a right to challenge reported performance with contrary evidence
 - NRAs should collate and publish incumbent service performance
 -
 - SLAs & SLGs should be tightly worded with limited carve-out conditions, clearly identify exceptions, limited opportunities for stop-the-clock, outage time for electronic platforms and other multiple feed-back loops

The L2 WAP product should also have unrestricted use i.e. it should be able to be used for consumer, enterprise or mobile backhaul service delivery.



3. Introduction

Vodafone has become a significant player in the fixed broadband market in a number of different countries, rising to become the largest Local Loop Unbundling (LLU) operator in Europe and winner of FTTH Service Provider of the Year 2014 (FTTH Council, Europe). Vodafone's LLU broadband is based predominantly on Central Office (CO) based ADSL2plus and SHDSL technologies. As broadband technology has evolved towards Next Generation Access (NGA), Vodafone's network investments have increasingly focussed on FTTH (using mainly GPON technology) and FTTC (using VDSL). In addition, Vodafone has acquired cable network assets in Germany, Spain and New Zealand. Vodafone is also a significant player in the business market. Key technologies classified as NGA are illustrated below:

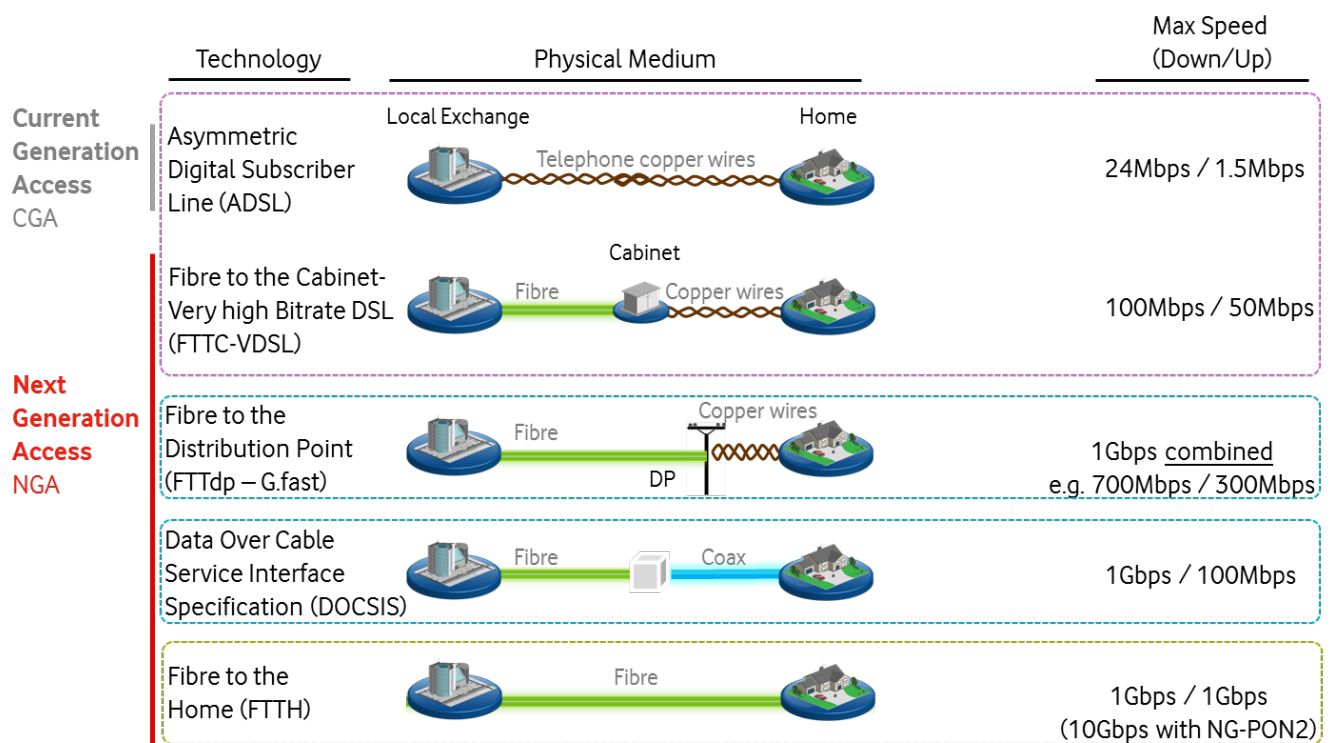


Figure 1:NGA Technologies

Despite significant investments, in any given local market Vodafone is unlikely to achieve 100% NGA coverage of the country by either build or buy approaches. Hence there will always be a need to use 3rd-party wholesale access providers (often, but not restricted to, the incumbent operators) for “off-net” access to end-user customers.

The preferred approach for off-net access is to use L2 WAP based on layer 2 Ethernet connectivity between the customer's premises and the hand-over point to the Vodafone network. L2 WAP is normally a regulated product; however the characteristics can vary significantly across local markets



and even between wholesale providers operating within the same local market. The EC is endeavouring to harmonise L2 WAP products across member states.

4. Architecture & Technology

4.1 Terminology

There are three business entities involved in the provision of L2 WAPs. These are explained below:

The **L2 WAP Network Provider** (or L2 WAP infrastructure provider) is responsible for the provision of the active and passive infrastructure over which L2 WAP¹ is delivered. The L2 WAP Network Provider offers standardised interfaces to which the ANO can connect, and delivers the L2 WAP user's traffic between these interfaces, across the L2 WAP domain. The L2 WAP domain extends from the end user premises to an interconnect point further up the network. L2 WAP Network Providers may own or lease the passive and active parts of the network, e.g. an L2 WAP provider may own the active electronics, but lease the passive infrastructure. This role is referred to as a Wholesale Provider by the Broadband Forum [1], but as a Layer 2 Wholesale Access Provider (L2 WAP) by BEREC [2] and as an ALA (Active Line Access) Provider by NICC [3]. This document will use the BEREC terminology.

The **Alternative Network Operator (ANO)** purchases Ethernet transport to an end user from the L2 WAP Network Provider over which it delivers services such as voice, video and internet connectivity. The ANO has a direct, contractual relationship with the L2 WAP Network Provider. The ANO may also have a direct relationship with the end-user or with other communications providers on a wholesale basis. ANOs may include ISPs and triple-play operators. This role is referred to as a Network Service Provider (NSP) by the Broadband Forum, as Alternative Network Operators (ANOs) by BEREC and as an ALA User by NICC.

The **End User** is the ultimate recipient of services provided over L2 WAP. End users include both residential consumers, and business users. They are sometimes referred to simply as customers or subscribers.

¹ The term "VULA" (Virtual Unbundled Local Access) is also often used in the context of this kind of L2 product. Strictly speaking, VULA only refers to the form of the layer 2 Ethernet wholesale service with local network interconnect/hand-over. This is then offered as an alternative to physical layer unbundling. Regional handover has more synergy with Bitstream in that it includes backhaul/aggregation network transport. In some countries (e.g Ireland or Italy) new Bitstream product variants have been introduced, often referred to as NGA Bitstream, which have the same VULA products characteristics apart from the handover location which is centralised instead of distributed.



4.2 Architectural Standards Context & Reference Models

The early DSL broadband architectures were based around ATM as the Layer 2 (link layer) technology. However, since around 2006 virtually all broadband networks have been based around Ethernet technology for layer 2, particularly for the aggregation network which backhauls access traffic to the core network and service nodes. A standardised approach for such architectures was developed by the Broadband Forum as Technical Recommendation TR-101[4]. As broadband network operators moved from CO-based DSL to NGA wireline technologies such as VDSL and GPON most leveraged their existing Ethernet broadband architecture that was developed for ADSL2plus and SHDSL access technology. The end-to-end IT systems, operational processes and broadband product propositions could be simply scaled to accommodate faster access technologies. Hence the approach preferred by broadband network operators at the forefront of NGA has been to simply replace or augment ADSL2plus access with NGA access within their existing Ethernet-centric broadband network architecture [5]. Adapting the Broadband Forum TR-101 architecture to include FTTH GPON access was standardised in Broadband Forum TR-156 [6] (illustrated below) :

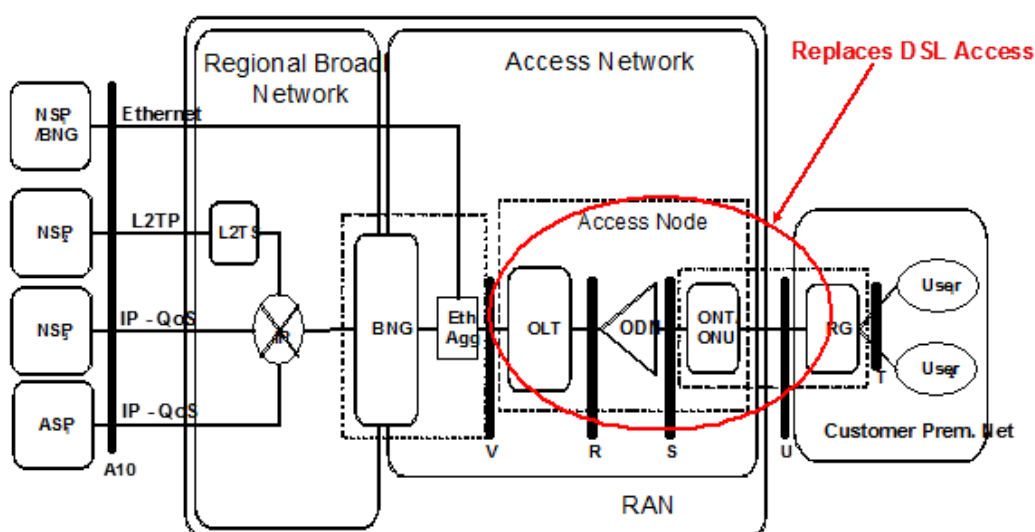


Figure 2: Broadband Forum TR-156 Reference Architecture for GPON

In the reference architecture above, the Ethernet interface to the “NSP” (Network Service Provider) is the NNI reference point applicable to a L2 WAP wholesale product. This enables the ANO to provide their own IP layer functionality and to provide their own BNG equipment for subscriber management purposes such as IP address allocation, user authentication and policy control (including lawful intercept).

As Ethernet-based Broadband network architectures became the norm for any network built after 2006, Broadband Forum standards work progressed to encompass all the requirements of multi-service broadband networks (see TR-144 [7]). These included the wholesale requirements necessary for L2 WAP. L2 WAP requirements were also well specified in [3] and [8]. These requirements were met by the architecture specified in TR-145 [9]. TR-145 extends the previous TR-101 and TR-156 Ethernet



architectures with the new technical requirements needed to fulfil the business requirements laid out in TR-144.

TR-145 “can allow network providers to offer Open/Equal Access Network involving a horizontal model comprised of Infrastructure Providers, Network Operators, Content Brokers, and Service Providers”. The scope of the TR-145 specifically includes the Ethernet Service Layer. This covers services seen by end-users as Ethernet Services (MEF Services) as well as the (emulated) Ethernet constructs to connect IP RGs or CPE to their IP Service Edges (such as a BNG). The Ethernet Service Layer performs 802.1ad Ethernet Aggregation, very much like the model deployed for in TR-101 architectures. An important piece of the TR-145 architecture is the concept of an Infrastructure Virtual Circuit (IVC) that is the basic building block for constructing end to end Layer 2 service connectivity as required by L2 WAP. An Infrastructure Virtual Circuit (IVC) can be from user (UNI) to service edge (NNI) as illustrated below:

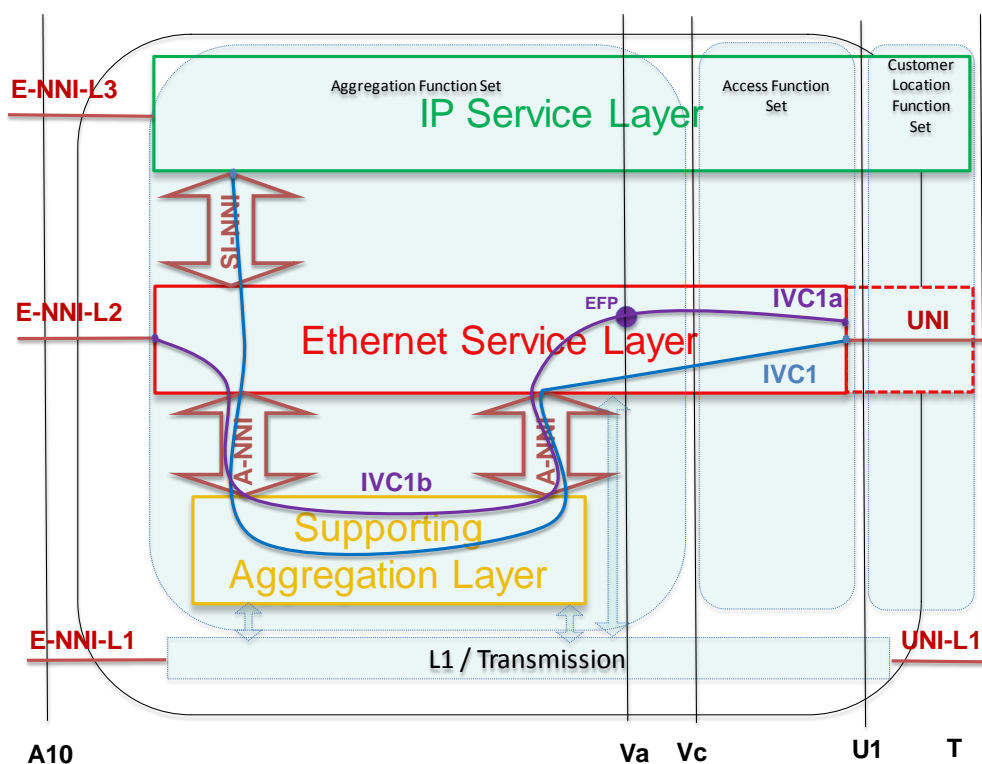


Figure 3: Illustration of 'L2 WAP-like' Ethernet Connectivity in TR-145 Reference Architecture

The L2 WAP is provided over this architecture by the Ethernet Service Layer between the U1 and A10 reference points. Other potential L2 WAP traffic handoff points also exist from the T reference point all the way to the A10 reference point. For example, traffic is handed off across U1 to the access nodes via access-specific transmission media (the so called “wires-only” interface for CPE). From the access nodes, traffic is handed off across Va (which could for example be in a CO) onto the Ethernet / MPLS aggregation network. The L2 WAP network must be able to deliver this Ethernet Service Layer independent of any underlying aggregation and tunnelling technologies.

The nodal (i.e. equipment) requirements to implement the L2 WAP-capable architecture are standardised in Broadband Forum TR-178 [1] (which also includes capabilities from [10]). This enables vendors to build equipment to meet the network architecture (TR-145) and business requirements (TR-



144) which include L2 WAP. The L2 WAP-specific standard reference model in TR-178 is illustrated below:




Figure 4: L2 WAP Reference Model from BBF TR-178 Standard

L2 WAP (as enabled by the Ethernet Service Layer in TR-178) is also known as the “L2 NSP Wholesale Model”. This enables L2 WAP networks to provide connectivity between both residential and business consumers (end-users) and their respective Network Service Providers (NSPs) in an open and flexible way. It uses Ethernet transport to allow an access network provider to offer logically unbundled access. The end user buys services from one or more NSPs (ANOs) who in turn buy service from the L2 WAP provider serving the end user. An Infrastructure Virtual Connection (IVC) is defined at the Ethernet Service Layer between the UNI (U or U1) and E-NNI-L2 (A10) interfaces. A10 is defined at the Ethernet layer allowing the ANO maximum freedom in how they wish to build their service by selecting their interconnect locations. In this way it differs from “Bitstream” wholesale broadband solutions which operate using PPP and L2TP, or IP, which generally requires centralised interconnect.

The context and relationship between the various key Broadband Forum network standards documents cited above for L2 WAP is illustrated below:

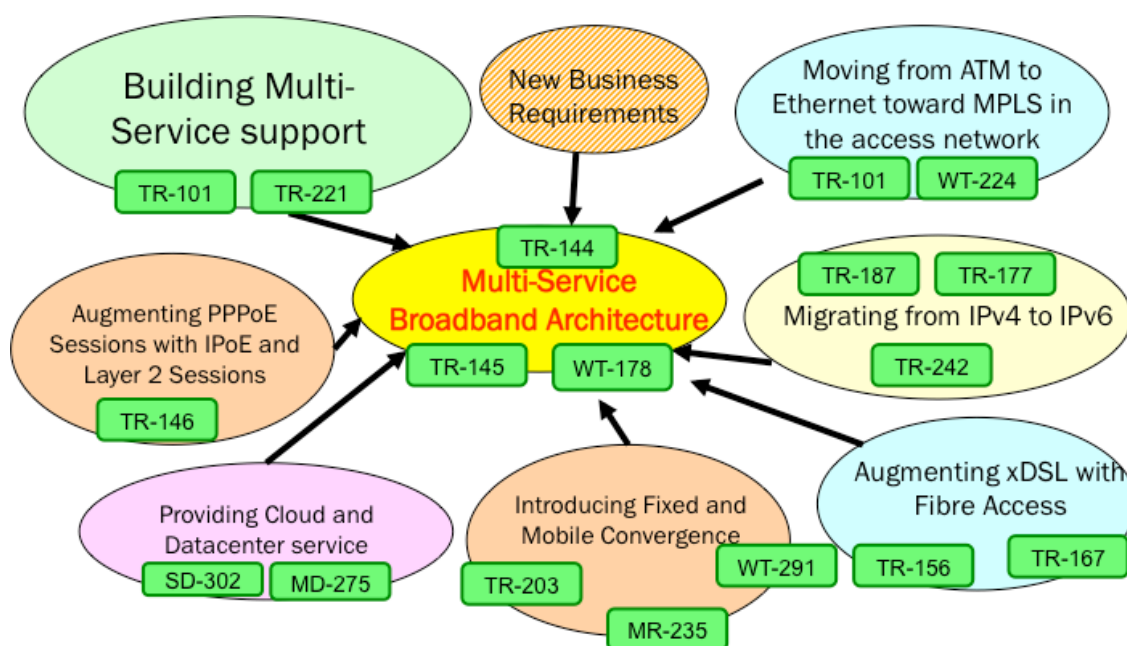


Figure 5: Context of Key BBF Network Architecture Standards Documents Relevant to L2 WAP

4.3 Ethernet VLAN Connectivity Options for L2 WAP

The L2 WAP is defined around the Infrastructure Virtual Connection (IVC) which provides an Ethernet connection between the ANO's network equipment and the CPE device on the end customer site that is generally operated by the ANO or by the end customer. Specifically, this IVC forms a logical circuit between the External Network to Network Interface (E-NNI) between the ANO and the L2 WAP operator; and the User to Network Interface (UNI) between the L2 WAP operator and the end CPE device. The IVC should be capable of transporting Ethernet MTU frames with a maximum size of at least 1600 bytes (excluding pre-amble and inter-frame gap). Support for mini-jumbo frames of 2000 bytes is desirable.

The L2 WAP can be built out of two components – access and backhaul [11]. Typically, an end-user is served by a single L2 WAP Network Provider. If an ANO does not interconnect in a distributed manner with the L2 WAP Network Provider, the ANO may extend the L2 WAP L2 Ethernet wholesale service to another handover location by connecting through a Backhaul Provider Network. The transport technology used within the L2 WAP/Access Provider Network and Backhaul Provider Network could be based on Provider Ethernet or MPLS, but the handover is Ethernet. In the L2 WAP network, an E-NNI for interconnect with ANOs (see Figure 4) could be supported directly by the Access Node or alternatively, a L2 Aggregation node (i.e. Ethernet switch) could be used to provide the E-NNI.

At the E-NNI

It should be noted that the NNI location with the lowest cost per IVC will be directly off the Access Node, as this prevents the costs of the aggregating node or MPLS network being loaded onto the cost of the services. However this requires a distributed network build or to separately purchase a backhaul service.



1Gbps Ethernet is considered the minimum interconnect speed, but options for 10Gbit/s Ethernet, and for any E-NNI to be protected using Link Aggregation (LAG) along with the Link Aggregation Control Protocol (LACP) should be available. This is especially important in order to minimise the chances of an NNI Port failure resulting in a major outage. The details for NNI protection are defined in [12]. Furthermore, the option to provide this protection from multiple Nodes (so called multi-chassis LAG) is preferred by the L2 WAP provider wherever possible.

At the E-NNI, the IVC is either identified by a specified S-VLAN ID or and S-VLAN ID and C-VLAN ID tuple. For the N:1 mode of operation as defined in standards, the end-user MAC address is also included in the tuple that identifies the IVC.

At the UNI

The network presentation at the end-user premises is either Ethernet (100 Mbit/s or 1 Gbit/s) or 'wires-only' (see more detail later in section 5.2). The customer's connections (IVCs) can be identified either by port or a port and VLAN ID tuple. At the U Reference Point there is usually a single physical interface per end user. In order to support multiple IVCs at the U Reference Point, these IVCs must be identified using VLAN tags.

VLAN Model Overview

A common example of a VLAN connectivity and forwarding model for NGA networks is illustrated in Figure 6.




Figure 6: Example of a Common VLAN Architecture for NGA

For residential delivery, multiple VLANs are offered on the UNI, allowing the CPE to deliver multiple separated services simultaneously. This is generally expected to be up to four services covering unicast data, multicast video, unicast voice and potentially VPN services. The VLAN numbering is expected to be the same on every customer user interface to ensure maximum re-use of the configuration.

On the uplink from the Access Node, the outer VLAN is also used to identify the service into which all residential customer traffic of that service is aggregated (multiple customers are aggregated into the same SVLAN). There are then two main options for identification of the end-user:



Option 1 (Figure 7): A tuple consisting of the Service VLAN and the customer MAC address is used to identify the end-user at the handoff. This is an N:1 model as defined by the broadband forum. In this case, many end-users share the same VLAN and additional security extensions defined by the Broadband Forum in TR-101 [4] are required to be implemented to protect against spoofing and denial of service attacks. This is not preferred by Vodafone for unicast L2 WAP handoff as it leaves critical MAC level security functions that separate customers in the hands of the L2 WAP provider.

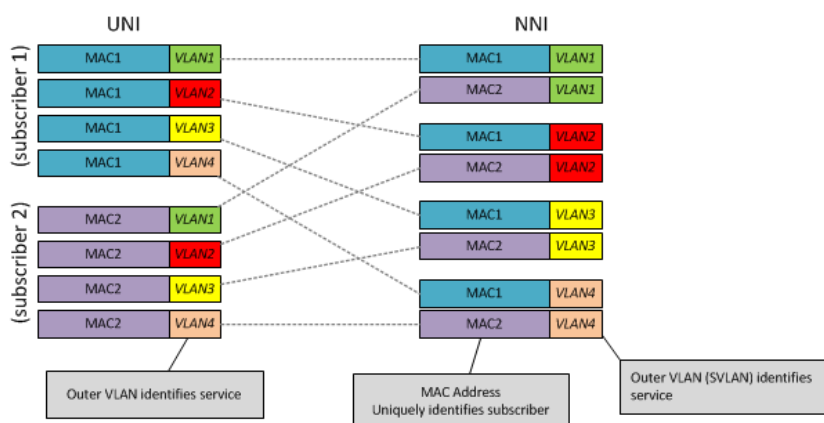


Figure 7 - Standards based N:1 VLAN mapping

Option 2 (Figure 8): A tuple consisting of an outer S-VLAN that represents the service and an inner C-VLAN that represents the customer is used to identify the end-user at the handoff. In this case, the S-VLAN is also an aggregation of multiple end customers, but the inner C-VLAN means that customers are separated from each other and can be uniquely identified. This is Vodafone's preferred L2 WAP model for unicast service hand-off because it matches our self-build architectures whilst also keeping the customer separation at the VLAN level.

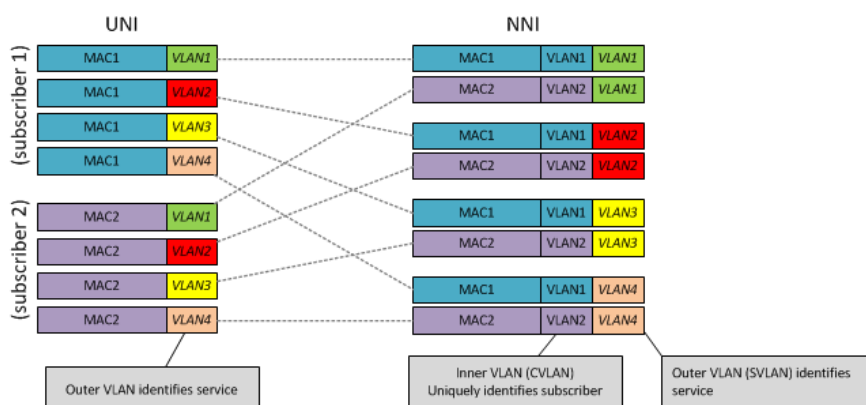


Figure 8 - SVLAN mapping with CVLAN per end-user (Vodafone unicast preference)

An inner VLAN on the uplink from the access node is then used to identify the consumer. This approach can increase the flexibility for rehomeing customers should a topology change be required. It also allows end-to-end monitoring to be done at an aggregate level. The E-NNI/SVLAN/CVLAN tuple then provides a unique identifier of the end-user and service.



For medium to large business customers, the preferred build requires a single SVLAN on the uplink from the Access Node to identify all traffic from the particular business end-user. The scalability is over 4000 S-VLANs (business end-users) per E-NNI, hence typically per CO. This is because the S-VLAN per service approach for residential end-users won't consume many VIDs in the S-VLAN address space (VIDs 1 to 4095). Any service VLAN on the business end-user's UNI is then mapped into this SVLAN (see Figure 9 below). The E-NNI/SVLAN/CVLAN tuple then provides a unique identifier of the business end-user and service. In this case, the VLANs on the UNI could vary based upon the enterprise service needs rather than being fixed as in the residential case.

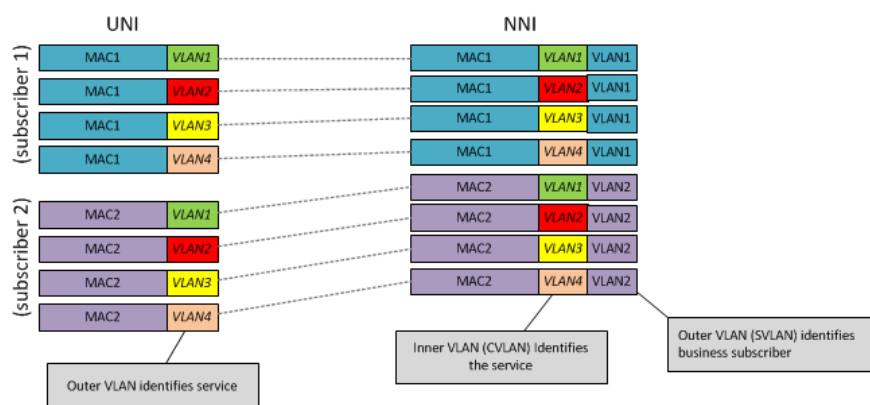


Figure 9 - 1:1 S-VLAN per business end-user

Service Edge Models:

It is desirable to have the flexibility to support both multiple-edge and single-edge service models, and the Vodafone preferred VLAN model above supports this.

The multiple-edge model has the following characteristics:

- End-user and service provisioning is in the access/aggregation networks.
- There is a shared VLAN per service for all unicast services, terminating in different service edge routers.
- There is a shared Multicast VLAN for IPTV broadcast services.
- The policy is per end-user/subscriber and the QoS management is distributed over access, aggregation and BNG.
- Access and aggregation is 'service-aware'.
- Optionally, a C-VLAN identifying the customer can be transported from CPE to service edge/BNG

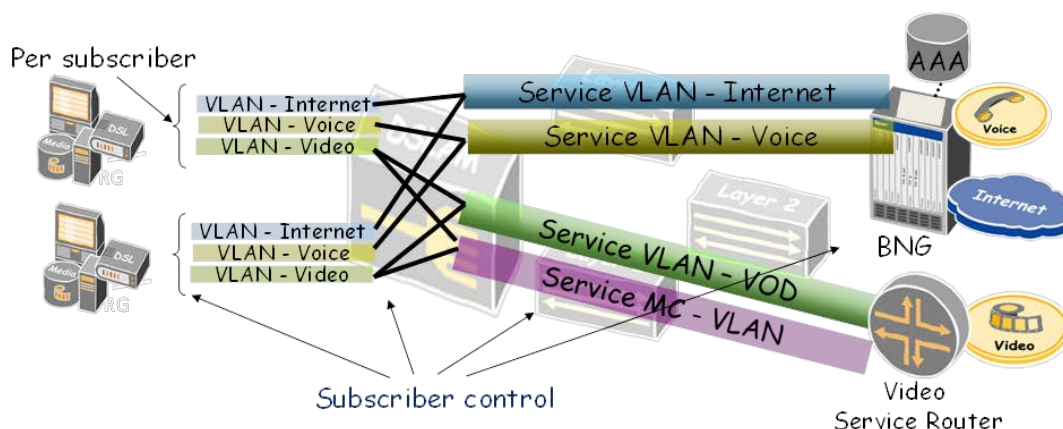


Figure 10: Multiple Edge Model (Service-Centric) for Residential Customers

The single edge model can also be implemented with the same architecture and has the same characteristics apart from the following:

- There is a shared VLAN per service for all unicast services, terminating in the same edge router (BNG).
- Optionally, a C-VLAN identifying the customer can be transported from CPE to service edge/BNG. In this case a subscriber centric approach is possible.

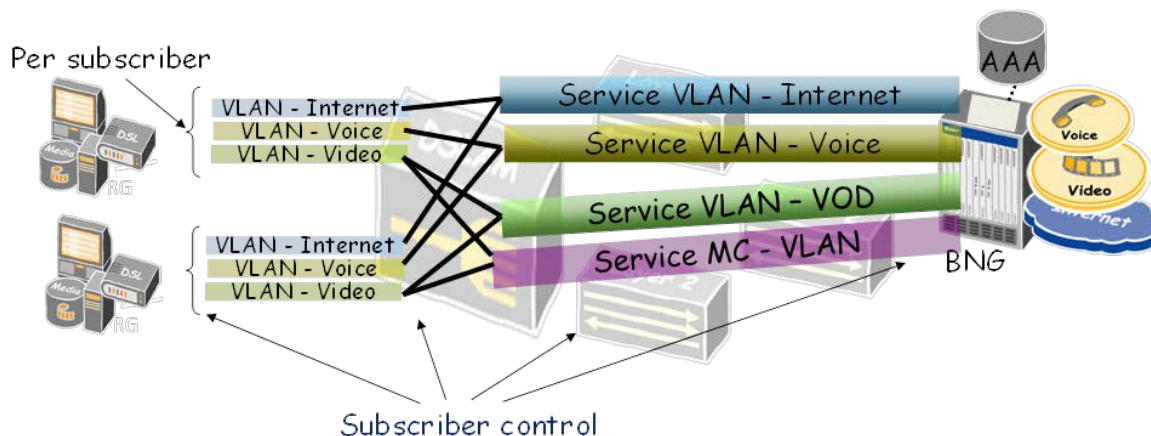


Figure 11: Single Edge Model (Service-Centric) for Residential Customers

The 1:1 model for business customers offers enhanced end-user/subscriber control and management and has the following characteristics:

- All end-users and services are provisioned in a single business service edge (such as IP-VPN PE router or BNG).
- There is an S-VLAN per end-user for all services.
- It is recommended to have a C-VLAN for each individual service.



- The policy is per-subscriber and the QoS management is centralized in the business service edge.
- Access and aggregation focus is on connectivity/transport.

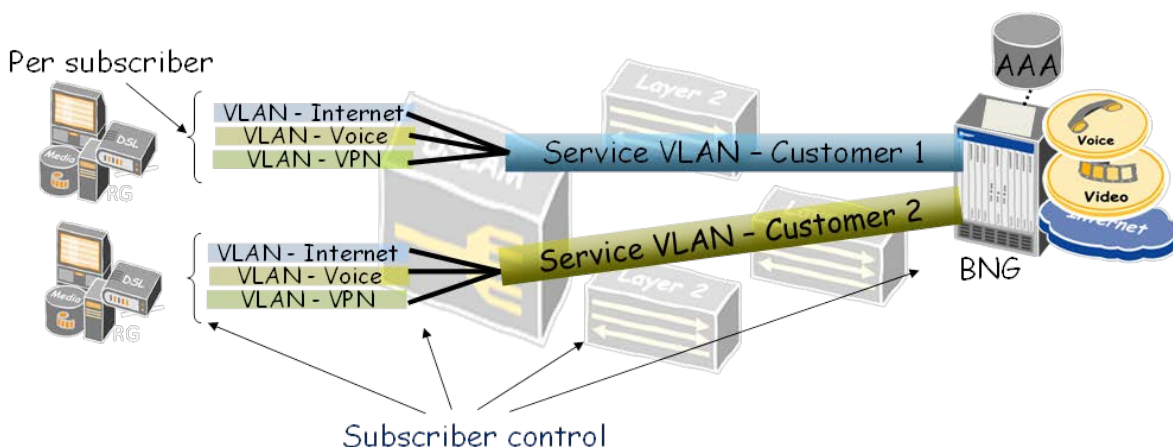


Figure 12: 1:1 VLAN Model for Business Services

All models use the same approach on the access line at the UNI, by defining a VLAN for each service. For example, as illustrated in Figure 6, separate VLANs can be used for:

- High Speed Internet (HSI)
- VoIP from CPE, if applicable
- Unicast/Multicast Video for residential customers
- VPN

NICC and BBF UNI and NNI endpoint mapping specifications

For a L2 WAP network product, the important aspects of the VLAN connectivity from the ANO's perspective are the "end-point maps" at the UNI and NNI, as these define how the customer traffic is mapped to the IVC. The internal workings of the L2 WAP network are a black box as far as the ANO is concerned. The network virtual connections need to meet performance, availability and security criteria but how exactly the network achieves this does not need to be explicitly prescribed.

The NICC and the Broadband forum identify a number of different end point mappings at the UNI and E-NNI:

A port based UNI: All frames received at the UNI irrespective of VLAN are mapped to a single IVC. Only a single class of service is supported. Any VLAN will be considered as user data from the perspective of the L2 WAP provider and will be exposed once again on the NNI without having been changed or stripped from the frame.



An S-tagged UNI: The outer TAG is used to identify the VID (i.e. VLAN Identifier, although a default can also be specified for untagged or priority tagged frames). Any inner VLAN will be considered as user data from the perspective of the L2 WAP provider and will be exposed once again on the NNI without having been changed or stripped from the frame.

A customer edge port based UNI: A single multicast VID is identified along with a corresponding VLAN ID. Any traffic with this VLAN ID is mapped to the multicast VID. All other frames are mapped to a default VID irrespective of VLAN and any VLAN will be considered as user data from the perspective of the L2 WAP provider and will be exposed once again on the NNI without having been changed or stripped from the frame.

Single Tagged NNI: In the single tagged NNI, the VID is identified only by the outer VLAN ID for a 1:1 architecture. The single tagged NNI also supports the N:1 architecture, but in this case, the end VID is identified by the combination of outer VLAN and MAC address.

Double Tagged NNI: In the double tagged NNI, the VID is uniquely identified by the tuple of both the outer VLAN and the inner VLAN. The standards specify that N:1 architectures are not supported in this model.

Mapping Vodafone Preferred model to L2 WAP standards [1] [10]:

As can be seen from the above description of endpoint mapping options, there is huge flexibility in the standards for the L2 WAP provider to offer different options. The following identifies the ability to meet the Vodafone preferred models with the Broadband Forum and NICC standard UNI and NNI combinations:

UNI mapping	NNI mapping	Residential Services	Business Services
Port based	Single Tag	Can support a single unicast service if the L2 WAP operator supports N:1 VID Does not map well to Vodafone preference. Only one class of service and difficult to optimise multicast within the L2 WAP operator domain.	Can support the Vodafone preference, provided that the VID is of the highest QoS required. Business service VLANs are transparently passed to the business service node. Would not require per service configuration in the L2 WAP network (This also maps well to the MEF UNI tunnel Access – MEF28)
Port Based	Double Tag	Will support a single unicast service, but does not map well to Vodafone multi-service preference. Only one class of service and difficult to optimise multicast within the L2 WAP operator domain.	Could support one business service only as otherwise we will see three VLAN tags on the NNI.
S-Tagged	Single Tag	Can support the Vodafone model, using standard N:1 model with MAC addresses to identify end subscribers	Can support Vodafone preference, although will require specific configuration in the L2 WAP operator



		(not preferred).	network for every business service VLAN to be added.
S-Tagged	Double Tag	Provided that full flexibility is granted to be able to specify S-Tags and C-Tags on the NNI and UNI, this can exactly replicate the Vodafone preferred multi-service model.	Can support Vodafone preference, although will require specific configuration in the L2 WAP operator network for any additional business service VLANs to be added,
Cust edge port	Single Tag	Supports a model where the Service is identified by the inner tag rather than the outer tag on the NNI (as it would be transparently passed). Can support the Vodafone model, but using N:1 model with MAC addresses to identify end subscribers and with VLAN re-mapping in the backhaul network (not preferred).	Can support the Vodafone preference with QoS. Business service VLANs are transparently passed to the business service node. Would not require per service configuration in the L2 WAP network (maps to the MEF UNI tunnel Access – MEF28, provided that no multicast VLAN is specified)
Cust edge port	Double Tag	Supports a model where the Service is identified by the inner tag rather than the outer tag on the NNI (as it would be transparently passed). Can support the Vodafone model, but using N:1 model with MAC addresses to identify end subscribers and with VLAN re-mapping in the backhaul network (not preferred). Would work well for a single unicast service solution.	Could support one business service only as otherwise we will see three VLAN tags on the NNI.

Table 1: Mapping of Vodafone VLAN Preferences to Standardised Options

Whilst it can be seen that no single standards-based option provides the perfect coverage of both residential and business services, an S-Tagged UNI with S+C tagged NNI would provide the closest solution for combined residential and business services, but a better solution uses a mixture of approaches on the same NNI in order to best support business and residential. This is already supported by some L2 WAP providers today and allows both business and residential services to use the same interconnect.

It is critical in all cases that Vodafone is able to have full flexibility to define the VLAN numbering approach on the NNI and UNI, as otherwise the NNI implementation may not meet our requirements to aggregate the customer services effectively such that the outer tag on the NNI identifies the service at the residential UNI. This will generally require a VLAN cross-connect capability in the L2 WAP Network Provider's equipment.



NETWORK ARCHITECTURE AND VLAN CONFIGURATION BEST PRACTICE

- ✓ **Support for S-Tagged UNI with a VLAN per service at the residential customer UNI**
- ✓ **Support for S-Tagged or Port Based UNI for business customer UNI.**
- ✓ **Support for S-Tagged or S+C Tagged on the same NNI.**
- ✓ **Minimum of 4 C-VLANs per residential customer with support for the same set of VLAN IDs on all UNI**
- ✓ **Full flexibility for the access provider to define the ranges and policy for VLAN allocation on the UNI and NNI.**
- ✓ **Access to standards N:1 VLAN for multicasting**
- ✓ **Unicast services can be identified uniquely using only VLAN IDs (not MAC address)**
- ✓ **VLAN per customer at NNI and UNI for business customers**

Best Practise Key Points 1: VLAN Configuration

5. Network Characteristics

5.1 Network Interconnect

A flexible L2 WAP wholesale product will offer ANOs a range of options for how and where they interconnect to the L2 WAP Provider in order to collect the traffic from their broadband end-users who are connected via NGA. In common with many other products that involve some form of “interconnect”, it is possible to conceive of at least three product options: National, Regional and Local (the technical interface requirements are well specified in [11] and [9]).

With a National variant of the product, the ANO would be procuring backhaul and core bandwidth from the L2 WAP infrastructure provider who would use their own network to transport the aggregated NGA traffic (from all NGA connections anywhere in the country) to the interconnect location² selected by the ANO, usually at one of the ANO’s major Points of Presence (PoP) or a 3rd-party “Carrier Hotel”.

A Regional product variant would enable more distributed interconnection points at a number of regional metronodes which act as aggregation points for all NGA connections (“access tails”) within a regional geographic area. This enables the ANO to leverage their own core network capacity (and hence this Regional L2 WAP product should be cheaper than the National interconnect variant) but the ANO is still using backhaul aggregation network capacity from the L2 WAP infrastructure provider.

The Local variant of the L2 WAP product goes a step further and enables the ANO to collect the NGA traffic directly at the local CO where, for example, the GPON OLT (and perhaps an adjunct Ethernet switch) are located. This enables the ANO to use their own backhaul or “middle mile” aggregation network capacity or to procure this from a 3rd-party who is not the L2 WAP infrastructure provider. This Local L2 WAP interconnect product option (which should be the cheapest) will be of particular interest

² There would usually be at least two national interconnection points to provide resilience



to LLU operators who could (technically) then leverage their existing LLU space, power and fibre backhaul connectivity. The interconnect is very simple. With the appropriate product options for the Ethernet VLAN stacking and numbering on the L2 WAP interconnect, the integration of L2 WAP by an LLU operator can be almost seamless.

As noted previously, robust interconnect arrangements should include the option of having local interface resilience between Vodafone equipment and the CO. Even if not taken up initially, the ability should be sought to easily request an upgrade to add another interface port into a Link Aggregation Group or similar mechanism.

Initial sizing of the interconnect needs to reflect plans for the service mix and also the ease with which it can be increased – multicast requires a steady allocation of bandwidth on the interconnect. LAG from more than one node is an option to improve resilience. 10G (or resilient via 10G+10G LAG) may be best from day one, since it future proofs the link. However, costs may dictate that 1G or n*1G be used instead, then provision additional 1G interfaces as demand grows.

Vodafone has core and backhaul (metro aggregation) networks in the majority of the local markets in which it operates. These provide network infrastructure to deliver mobile and/or fixed network services. It is therefore Vodafone's preference to leverage its existing network infrastructure as much as possible and to only use third-party networks where Vodafone doesn't have a network presence. This should reduce the cost of using the third-party network (since 'less' of it is required) and gives Vodafone a greater degree of end to end control. In the context of using L2 WAP, the preference is to restrict its use to "last mile" access and hence to interconnect to the wholesale L2 WAP provider as close as economically feasible to the end customer. This often means having a Network to Network Interconnect (NNI) within the local CO. The diagram below illustrates the prioritised list of network interconnect locations (which includes both L2 WAP and bitstream scenarios for greater context).

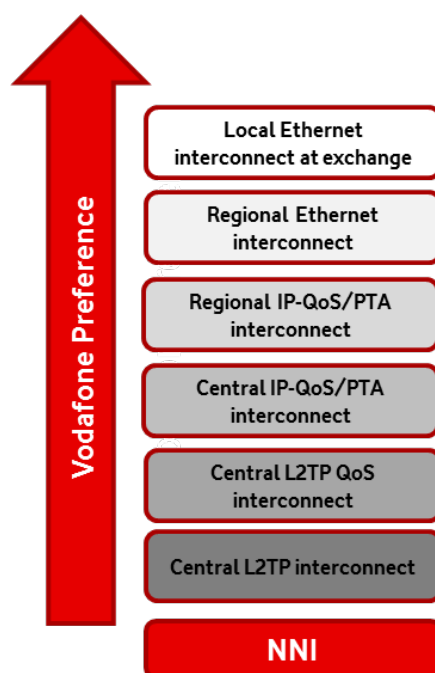


Figure 13: Prioritised List of L2 WAP & Bitstream NNI Locations and Approaches



It is unlikely that Vodafone has a presence within every CO from which an incumbent offers NGA. Indeed, providing our own backhaul fibre to every CO is likely to be uneconomic. However, having interconnect in just a few regional aggregation nodes does not maximise use of our own backhaul/aggregation network and so is also economically sub-optimal. An example of a best practise approach is the “parent-child” CO arrangement. In this scenario, the incumbent wholesale NGA provider aggregates the NGA traffic from a few (typically 5) “child” COs to an NNI location at a “parent” CO. Hence, build-out of our own fibre backhaul to each Parent site gives coverage of all NGA customers in its Children COs. This is illustrated below where the lines show the Parent to Child CO links and coloured areas show the aggregation of the CO areas. This example of best practise is available from BT Openreach in the UK where an ANO only has to build fibre backhaul to an NNI presence in ~1200 Parent COs in order to access NGA customers connected to ~5000 Child COs. Conversely, where economically and technically feasible, there could be an option to hand over at a lower level than the CO (e.g. cabinet in an FTTC scenario).




Figure 14: Illustration of Parent-Child CO Approach for NGA Aggregation to NNIs



L2 WAP NETWORK INTERCONNECT BEST PRACTISE

- ✓ **Ethernet Interconnect at CO Level**
- ✓ **Parent-Child CO aggregation hand-over option and the possibility to request lower level hand-over at a street cabinet**
- ✓ **1G, 10G and resilient handover options (inc. n*1G & n*10G)**

Best Practise Key Points 2: L2 WAP Network Interconnect

5.2 CPE/Modem Choice

Customer Premises Equipment (CPE) has proven to be a key domain for service differentiation and branding in broadband access deployments. The broadband CPE market follows a standard evolution trajectory:

Phase 1: Initially the L2 WAP Provider will supply and install the CPE to ensure it is interoperable with their network equipment and installed correctly to avoid provisioning failures. This “CPE” may actually be just a modem or NTE. Often the supplied NTE is produced by the same vendor as the network equipment. The ANO may be permitted to supply a separate L3 router. Hence it can result in a 2-box solution.

Phase 2: The L2 WAP Provider may then offer a limited range of CPE, perhaps from two or three different vendors to give the end-user some choice over the combination of price and functionality. This CPE will have been thoroughly pre-tested by the L2 WAP Provider. Over time, this approved list (sometimes called a “whitelist”) of CPE will grow³. ANOs may be offered the chance to have CPE from their own preferred vendor “qualified” (via lab testing) for use with the broadband access connectivity product and hence added to the approved CPE list. A self-install option may also be introduced at this stage once the L2 WAP Provider has digested the lessons learned from its own installation experiences and produced best practice installation guidelines together with associated trouble-shooting flow-charts and support tools.

³ For DSL technologies, an alternative approach to a CPE whitelist could be a model which relies on the interoperability of the major xDSL chipsets in reference implementations, for example by using the latest two official driver versions. VULA Network Providers would need to maintain interoperability for the DSLAM models they use and for each DSLAM release upgrade. ANOs would require that their CPE vendors strictly implement the chipset reference designs and use unaltered DSL drivers. This could initially be proven by comparing the performance of the actual CPE implementation with the chipset reference design.



Phase 3: Once equipment standards and vendor interoperability have sufficiently matured (often following numerous industry “plugfest” events) and self-installation processes have proven themselves to be sufficiently robust, the L2 WAP Provider may be ready to move to a full retail model. In this final phase an end user can procure CPE from a variety of retail outlets (both high-street stores and on-line) to meet the price/functionality point of their preference. This final phase also effectively makes the Access connectivity product a “wires-only” product (since no NTE is provided by the L2 WAP Provider).

ADSL2plus and SHDSL are just two examples of technologies that have already followed this CPE market evolution path through to phase 3. VDSL is currently at Phase 2 with some L2 WAP Providers whereas GPON is starting to move from Phase 1 to Phase 2 (the interoperability work of the Broadband Forum has greatly helped prepare the technology for the move to Phase 2).

There is no good reason why GPON CPE can't soon reach the retail model of Phase 3. This is exemplified in Vodafone Portugal where successful interoperability between GPON OLTs and ONTs from different vendors has been proven. It will take commitment from regulators (to advocate it on L2 WAP roadmaps) and work in the standards bodies and interoperability test labs but this objective is achievable. It wasn't long ago that skeptics were saying that a self-install retail model for DSL was impossible, but now we have a thriving and competitive DSL CPE industry in many local markets. It is important to establish the vision for the “strategic end game” so that the GPON industry can chart the appropriate course with respect to product development roadmaps, standards and interoperability events.

In order to reach the mature GPON CPE market scenario whereby the ONU/ONT is provided by the ANO or end-user (as opposed to the L2 WAP Infrastructure Provider) capabilities are required in three key areas:

1. Integrated CPE that avoids the space, power and home-wiring installation complexity required by multi-box solutions.
2. A standardized approach to virtualising the management of the CPE (especially in integrated solutions) so that CPs and the L2 WAP Infrastructure Provider can take responsibility for assuring the end-end service and network components that they are responsible for.
3. Standardised OAM features spanning the GPON physical layer up to the Ethernet and IP layers that are key for assurance and hence fault demarcation for anything that can go wrong (particularly during installation).

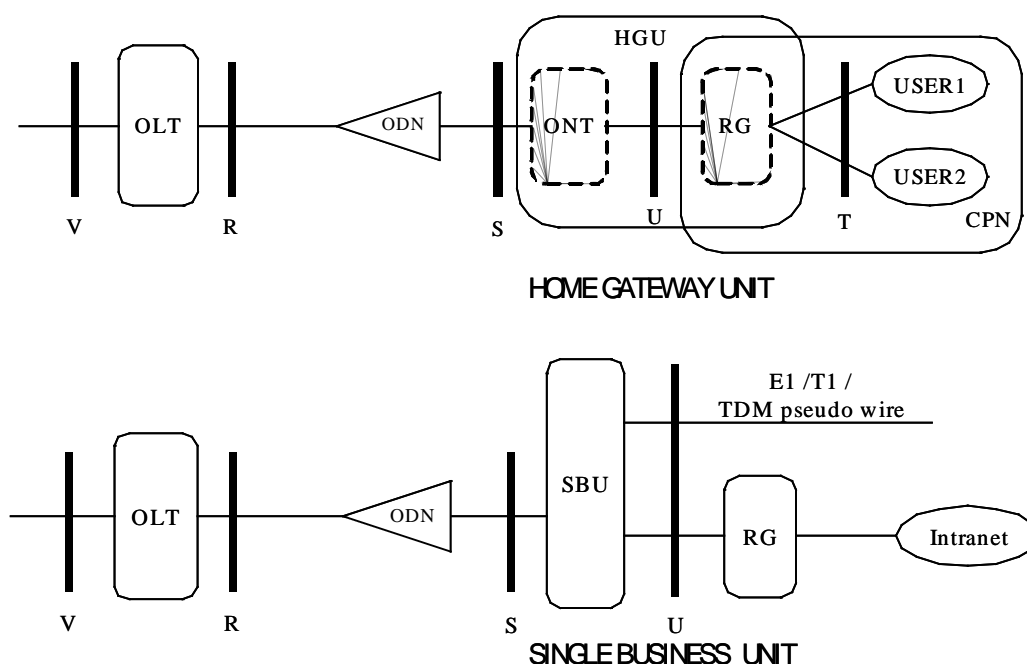


Figure 15: Two Examples of Integrated GPON CPE Designed to Meet Different Market Requirements

The figure above illustrates that there isn't a satisfactory "one size fits all" solution when it comes to integrated GPON CPE. In addition to the above examples, others can be conceived for different market segments such as those optimised for multi-tenancy units of different sizes. Indeed, many GPON equipment vendors already produce a coherent "family" of GPON CPE products tailored to the various target markets. Some are also looking to introduce an ONT in the form-factor of an SFP that can plug into a standard SFP cage.

The importance of integrated CPE is further illustrated by examining the delivery of a triple-play service bundle (including broadcast video) to consumers. The functional elements required to deliver such a service bundle over GPON can include :

1. ONT (to terminate the GPON physical layer transmission)
2. Battery Back-Up Unit (to keep services alive if mains power fails)
3. Router (for data/Internet access)
4. Multicast functionality and IGMP support (to deliver IPTV⁴ video to Set Top Boxes)

⁴ The focus is on IPTV as opposed to the less integrated IRS overlay approach which is a short-term expedient and not compatible with the longer term objective of delivering multiple integrated services over a single IP/Ethernet network



- 5. Analogue Terminal Adaptor - ATA
(to present voice services to existing analogue phones)
- 6. Set Top Box

These six functions could be physically provided in five different boxes (assuming the multicast functionality resides in the router). However, this would not be elegant in terms of space in the home or the complexity of all the interconnecting cabling to install (and debug when there is a fault). Wiring closets and cabinets can tidy-up the aesthetics and cabling but essentially this is just tantamount to placing five smaller boxes in one bigger box. Such an approach is also against the spirit of the European Commission code of conduct on reducing power consumption in broadband equipment [13]. If the Set-Top Box is kept separate (as is normal in most triple-play deployments today) then the remaining five functions can easily be integrated into a single device. Indeed such equipment already exists from a number of vendors.

Beyond the CPE integration and interoperability for L2 WAP products we also need to “virtualize” the management of integrated CPE. This means providing a management architecture and associated B2B interface within the wholesale L2 WAP product that allows ANOs and the L2 WAP Infrastructure Provider to configure and query the functional elements and MIBs within an integrated CPE which relate to the service component that they are supplying. When the CPE is integrated, it is vital that each player can be restricted to accessing only their own service elements. A fundamental framework that can be leveraged for such a scenario for GPON is provided in Broadband Forum WT-142 [14]. This uses GPONs OMCI protocol for the L2 WAP Infrastructure Provider to configure the GPON physical layer and Ethernet link layer⁵ (setting up VLAN structures and IDs etc.). It then uses Broadband Forum TR-069 [15] and the associated family of related recommendations (e.g. [16][17]) to configure and report on the Internet, voice and video service elements within the CPE. The L2 WAP Infrastructure Provider could provide this product capability via a partitioned TR-069 Auto-Configuration Server (ACS). The principle is illustrated below :

⁵ ATM ILMI provided the equivalent functionality within early DSL CPE

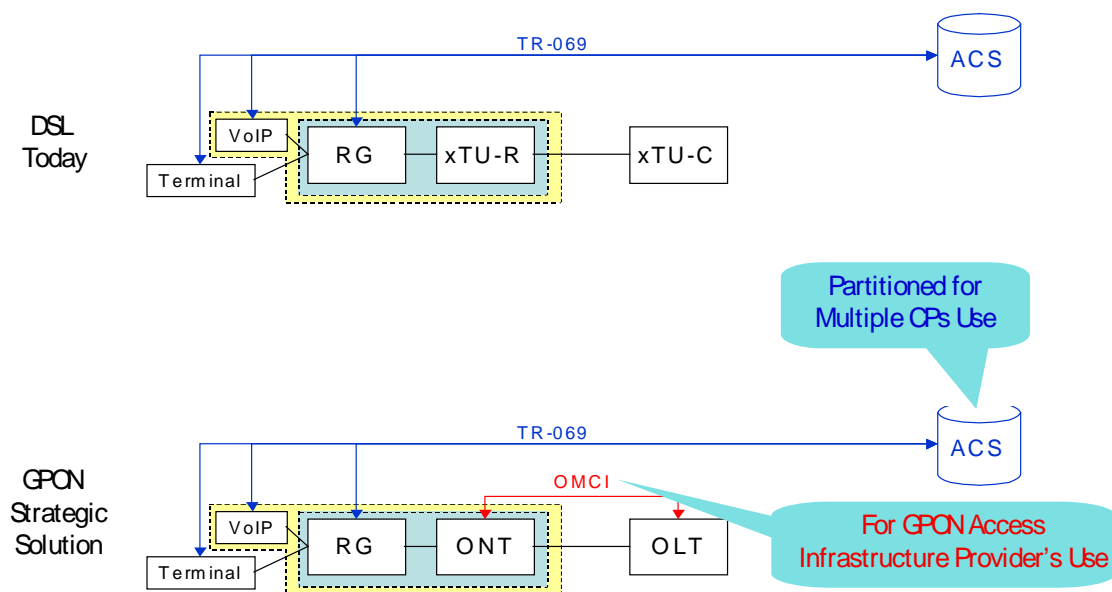


Figure 16: Partitioned GPON CPE Management with TR-069 Complimenting OMCI

In a situation where a single ANO was offering a bundled package of services such as triple-play they could provide their own ACS (leaving the L2 WAP Infrastructure Provider to just use OMCI). However, there may still be economies of scale for the L2 WAP Infrastructure Provider to offer the ACS as a service option that can be shared across all ANOs.

For GPON, ONTs are normally pre-provisioned on the OLT (using the serial number or password ONT activation method as described in TR-156 [6]). Hence for a “wires only” GPON CPE model, a joint process for ONT authentication will be required between the ANO and L2 WAP Network Provider.

In summary, the L2 WAP UNI (as per Figure 4) can support either an Ethernet or “Wires-Only” interface (the technical interface details are well specified in [18]).

- For Ethernet presentation, a L2 NT managed by the L2 WAP Network Provider will terminate the access loop technology and support a VLAN tagged, 802.3 UNI interface the U1 reference point. The L2 NT could incorporate a VDSL2 modem or GPON ONT. Each ANO can provide their own CPE that connects to the L2 NT.
- In the wires-only case, the L2 WAP Network Provider will provide a passive interface via a NID in the end-user premises. The end-user CPE will then terminate the access loop technology. This is referred to as a “wires-only” interface since the L2 WAP Network Provider’s product does not include an active Network Termination (i.e. a broadband modem). The Broadband Forum has developed a series of interoperability test specifications that cover both functional and performance tests [19], [20], [21]. These can be used to support certification programmes for VDSL, GPON etc. to facilitate wires-only UNI propositions for NGA L2 WAP.

Vodafone designs its own range of broadband Customer Premises Equipment (CPE) which includes a range of broadband CPE with a variety of WAN interfaces (Ethernet, ADSL, VDSL, GPON ...). This is usually



managed remotely using the Broadband Forum's TR-069 protocol. Hence, subject to UNI (User to Network Interface) interoperability between the CPE and network, Vodafone's preference is to provide its own broadband CPE. A prioritised list of preferences is illustrated below:

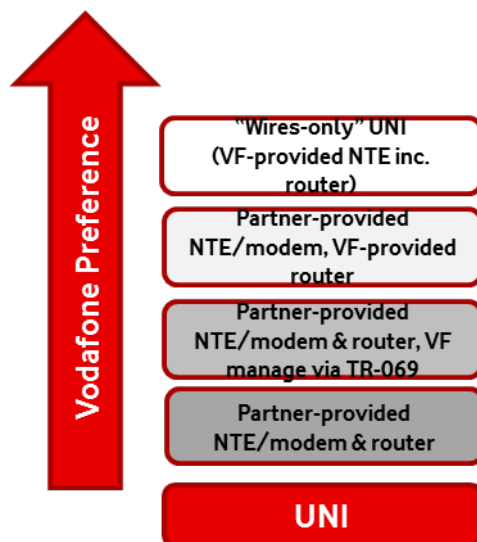


Figure 17: Prioritised List of UNI Preferences

Three user port types are feasible for L2 WAP:

- A port based UNI.
- An S-tagged UNI.
- A customer edge port based UNI

The S-Tagged UNI requires the ANO to identify the connection of a frame using either an SVLAN tag or the default S-VLAN on the port.

A customer edge port based UNI offers the same capability but in this case any C-VLAN tags used at the UNI are tunnelled over a point-to-point connection. This means that to use a VLAN tagged presentation at the UNI the ANO will need to send three VLAN tags at the NNI; where the two outer VLAN tags are defined by the virtual connection endpoint map (see [10]) at the NNI and the inner-most VLAN tag needs to match the VLAN tag being used by the ANO at the UNI. This innermost VLAN tag is not significant to the L2 WAP Network Provider at the NNI and is carried transparently over its network.

CPE BEST PRACTISE

- ✓ **Wires-only option for ANO to provide own branded & integrated NTE/modem/ONT + router CPE**
- ✓ **CPE interoperability requirements follow BBF standards and certification approach**
- ✓ **CPE interop test/validation environment to facilitate expanded "white list" of approved CPE**

Best Practise Key Points 3: CPE and UNI



5.3 Bandwidth, Traffic Prioritisation & Quality of Service

The bandwidth (both downstream and upstream speeds) of the end-user broadband access connection is a fundamental attribute of a L2 WAP network product. This dictates the type of services it is capable of supporting (e.g. Ultra-HD video). Some DSL-based NGA products⁶ can be operated in a rate-adaptive mode up to the maximum speed that the physical layer transmission system can deliver of the end-user's connection (rather than being constrained to just having tiered fixed-rates)⁷. In such circumstances the ANO should be able to access the maximum speed capability that is physically achievable for the end-user without being artificially constrained to a capped lower speed.

The ANO should be able to use maximum bandwidths at least equal to any offered by the retail service provider division of the L2 WAP Network Provider. For competitive reasons L2 WAP Network Providers should have to provide at least the capability for ANOs to access the same underlying profiles that are used in any retail products offered by the L2 WAP Network Provider. Ideally, ANOs should have equitable access to set the requirements for bandwidth/QoS levels, and when the design is agreed ANOs should have equal access to them compared to the downstream retail units of the L2 WAP Network Provider. L2 WAP is intended to replicate as much as possible physical access to enable innovation and service differentiation. QoS and bandwidth profiles are an essential part of that. A process to certify and add new bandwidth profiles (with SLAs on timeliness of making them available for deployment) is also desirable.

The L2 WAP network design and capacity management should ensure that a minimum guaranteed bandwidth is offered as part of the L2 WAP product. In other words transmission should be uncontended in practice and provide guaranteed bandwidths according to the access seekers' needs. The minimum throughput capability defined through a set of CIR profiles represents an uncontended bandwidth in practice. The minimum bandwidth should have an associated SLA. The L2 WAP product should also include an option where a larger proportion on the maximum available access bandwidth is uncontended so that it is suitable to be used for business services.

A L2 WAP network that meets ANO's requirements needs to support different services concurrently within the physical layer bandwidth constraints of the broadband access transmission system, therefore traffic prioritisation and QoS features need to be supported as part of the L2 WAP product functionality. However, QoS should only be invoked when network congestion occurs at aggregation nodes. The L2 WAP Network Provider should not rely on QoS alone but should combine it with operations processes for network capacity management and network demand (ingress load) management in order to minimise the potential for the ANO to experience Ethernet frame loss/delay on its services.

⁶ E.g. such as VDSL or G.fast. There are also a range of standard techniques (such as via use of DHCP & PPPoE IA) to communicate the DSL line's synch-rate to the ANO to enable them to shape traffic appropriately.

⁷ Bonding multiple DSL access lines (e.g. via EFM – Ethernet in the First Mile techniques) may also be used as a way to increase the physical layer speed offered.



The L2 WAP network architecture also needs to support multi service edge deployments. This means that downstream traffic from different sources (e.g. on different VLANs) could be converging on an Access Node to be conveyed over an end-user's broadband access connection. Examples of different traffic sources and service edges could include an IMS switch for voice, a multicast feed for live video and a BNG for Internet access. This convergence of traffic flows at a multiplexing node (such as OLT or VDSL cabinet MSAN) can cause potential congestion (even at sub-second levels) for individual end-user's traffic flows and in aggregate across multiple end-users sharing the same Access Node. Hence the traffic needs to be appropriately prioritised and scheduled to deliver appropriate Quality of Experience (QoE) outcomes. A minimum of four prioritisation levels is recommended, with Ethernet p-bits used to mark or colour the QoS class.

The L2 WAP network must be able to support different treatment for loss-tolerant services (can tolerate shaping, policing, oversubscription) and loss-intolerant services (bandwidth needs to be capacity planned). At network nodes where traffic is aggregated (such as MSAN and OLT Access Nodes), there is the potential for congestion at the aggregation interface. Hence there can be a need to adapt the aggregate speed of the Ethernet interfaces (at the A10, U1 and Va reference points in Figure 3) to rates lower than the access line-speed in order to interwork with underlying transmission infrastructure or to support hand-off traffic to L2 WAP customers at arbitrarily defined rates.

On ingress to the network the L2 WAP Network Provider needs to classify the Class of Service for each ingress Ethernet frame. These classes of service can have different frame delivery performance objectives. Each IVC can have associated bandwidth profiles for each Class of Service. To support hand-off to CPE at U and U1, the class of service mapping and policing of these bandwidth profiles in the upstream direction needs to be supported by a L2 NT and the Access Node. The L2 NT may need to schedule traffic into the Access Loop in the upstream direction in a way that is aware of the Classes of Service of multiple IVCs (e.g. VLANs). Similarly, in the downstream direction, the Access Node may need to manage any contention between multiple IVCs by scheduling traffic onto the Access Loop.

Note that the ANOs can also use their own Service Edges (such as a BNG) to implement per customer QoS in addition to the L2 WAP network capabilities.

The L2 WAP Ethernet Wholesale Service should support 4 QoS Classes (A, B, C, D). A, B and C have behaviour in common with the traditional EF, AF and BE Classes. Classes C and D are both BE traffic, but Class D can be constrained to a defined share of the BE traffic under congestion. While they are intended to support a variety of service types, some typical examples are given in the table below [1]:



QoS CLASS	TYPICAL USE
A	Real time delay sensitive applications (e.g. voice or packet-layer synchronisation)
B	Streaming applications (e.g. video)
C	Internet data
D	Guest or 3 rd -party access

Table 2: Example Uses for L2 WAP QoS Classes

Each of these classes should have an associated performance objective (i.e. Ethernet frame loss and delay) that form a per-class service level specification offered to the ANO. This service level specification or QoS SLA could be based upon the MEF 10.2 performance attributes:

- Frame Delay (percentile)
- Frame Delay (mean) *e.g. may range from say 10ms to 75ms depending on technology & QoS traffic class*
- Inter-Frame Delay Variation (percentile)
- Inter-Frame Delay Variation (range) *e.g. may range from say 2ms to 10ms depending on QoS traffic class*
- Frame Loss Ratio *e.g. may range from say 0.05% to 0.8% depending on QoS traffic class*
- Availability

The QoS SLA should quantify the performance objectives offered for each QoS traffic class. These performance objectives are such that Class A has absolute scheduling priority over Class B, which in turn will have absolute scheduling priority over Classes C and D. Note that the L2 WAP Network Provider may use an additional very high-priority class for “internal” network use (e.g. for control protocols such as those used for routing updates, failure detection or synchronisation purposes). However, this class is not available to the ANO to use.

Starvation of the lower priority queues can be avoided by the use of per Class Policers. Example implementations are described in Annex A of ND1644 [10]. Classes A and B support only committed bandwidth. The bandwidth available for each of Classes A and B may need to be restricted by the L2 WAP Network Provider to ensure that QoS performance guarantees can be delivered for lower priority classes and other services. Classes C and D support both committed and excess bandwidth. The bandwidth profiles at the UNI and NNI can be configured to be ‘colour’ aware so that the ANO’s drop precedence marking is respected within these classes. In the case that both classes C and D send excess traffic at the same time, the L2 WAP Network Provider should limit the bandwidth share of Class D. Typical use cases for Class D would be to support WiFi guest access at the end-user premises, or to limit the bandwidth of a background application such as push video.

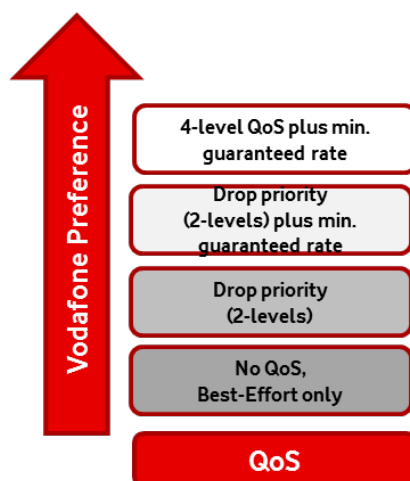


Figure 18: Prioritised List of QoS Preferences

In summary, the L2 WAP Network Provider should provide per VLAN QoS guarantees for the VLAN connections from the NNI to the UNI⁸ and also a CIR from NNI to UNI.

The L2 WAP Network Provider's product collateral should clearly explain to the ANO the QoS options and how traffic is treated. It should describe how traffic scheduling is performed (e.g. simple prioritisation or hierarchical). This may involve queuing, shaping, and policing to provide the proper treatment to traffic depending on its QoS classification, both on ingress and egress. Any traffic filtering mechanisms should also be explained. These could include conditional access control such as ACL filtering and multicast group control that enables the joining and leaving of multicast groups.

L2 WAP Bandwidth, Traffic Prioritisation & QoS BEST PRACTISE

- ✓ **Bandwidth profile available up to the maximum speed achievable by the DSL physical layer transmission system (including rate-adaptive systems)**
- ✓ **Ability to request new bandwidth profiles for FTTH transmission systems**
- ✓ **Choice of bandwidth profiles at least equal to those used by the L2 WAP Network Provider for their own retail services with the ability to request additional bandwidth profiles**
- ✓ **Minimum of 4 levels of traffic prioritisation based on p-bits**
- ✓ **Minimum quantified guaranteed throughput rates for upstream and downstream traffic (CIR from UNI to NNI) for over-booked and uncontended services**
- ✓ **QoS SLA targets for each VLAN from UNI to NNI, defined by Frame Loss Rate, Frame Delay and Frame Delay Variation measures**

Best Practise Key Points 4: Bandwidth, Traffic Prioritisation & QoS

⁸ Rather than just for shared QoS Class queues, which could result in competing for bandwidth with other ANOs that use the same VULA network.

5.4 Multicast

An obvious advantage of NGA over CO-based DSL is the ability to offer increased access speeds to end-users. Various service offerings can exploit this increased speed but one of the most often cited benefits of NGA over CO-based DSL is the ability to deliver multiple simultaneous HDTV channels. The concentration of local multicast service take-up has a direct impact on the cost-benefits of multicast network functionality. As we evolve to Ultra-High Definition (UHD) 4k TV there will be even greater bandwidth savings and economic benefits from using multicast so it is critical that such functionality is supported on wholesale L2 WAP products and architectures.

As an example to illustrate multicast functionality within an NGA transmission system let us consider GPON. GPON access systems are based upon point-to-multipoint transmission and so are inherently capable of supporting multicast (hence virtually all GPON vendor equipment includes multicast functionality). A single copy of a multicast video channel can be delivered to the GPON OLT (Optical Line Termination or “head-end”) which can then replicate the channel to all ONUs (Optical Network Units) connected to the OLT. Exposing this multicast functionality to ANOs as part of a wholesale GPON product enables ANOs to construct efficient end-end IPTV services and triple-play⁹ bundles.

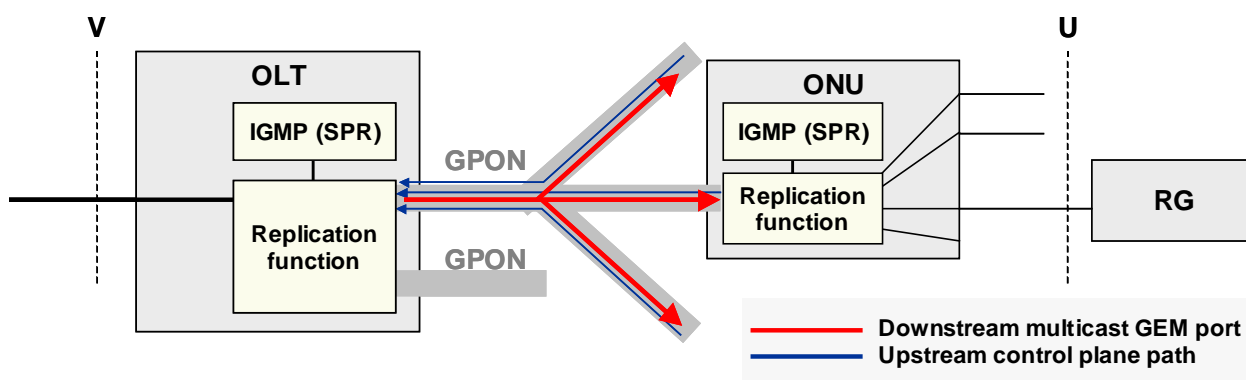


Figure 19: GPON Reference Model Showing Multicast Functionality

If the L2 WAP product includes the functionality to support multicast, then it will be feasible to deliver a single copy of a multicast channel to the Access Node (e.g. GPON OLT or VDSL MSAN) and have it replicated to all end-users. If this basic capability does not exist in the L2 WAP product offering then the ANO may need to deliver multiple copies of the channel to the Access Node (one for each customer wanting to watch it) and then the Access Node would use unicast techniques to deliver each of these copies to an individual end-user's CPE. This latter unicast approach is inefficient with respect to the end-end transmission path between multicast “head-end” and the end-user: Firstly, the backhaul network connecting the Access Node to the multicast head-end would need to carry multiple unicast

⁹ Triple-play = Internet access + Voice + Video where the video component virtually always means streaming “live” broadcast content and may optionally include video on demand



copies of the video channel instead of a single copy. Secondly, the NGA transmission system itself would need to carry multiple unicast versions of the video channel.

Multicast improves network efficiency by sending the same IP traffic to a group of receivers in a single transmission across a dynamically signalled multicast forwarding tree that is set up as a result of end-user signalling. Broadband Forum TR-101 [4] describes the usage of IP Multicast Routing on the BNG, and RG and IGMP snooping/proxy on the access nodes and L2 aggregation nodes to set up this forwarding tree for IPTV services. The TR-101 multicast reference model is reproduced below:

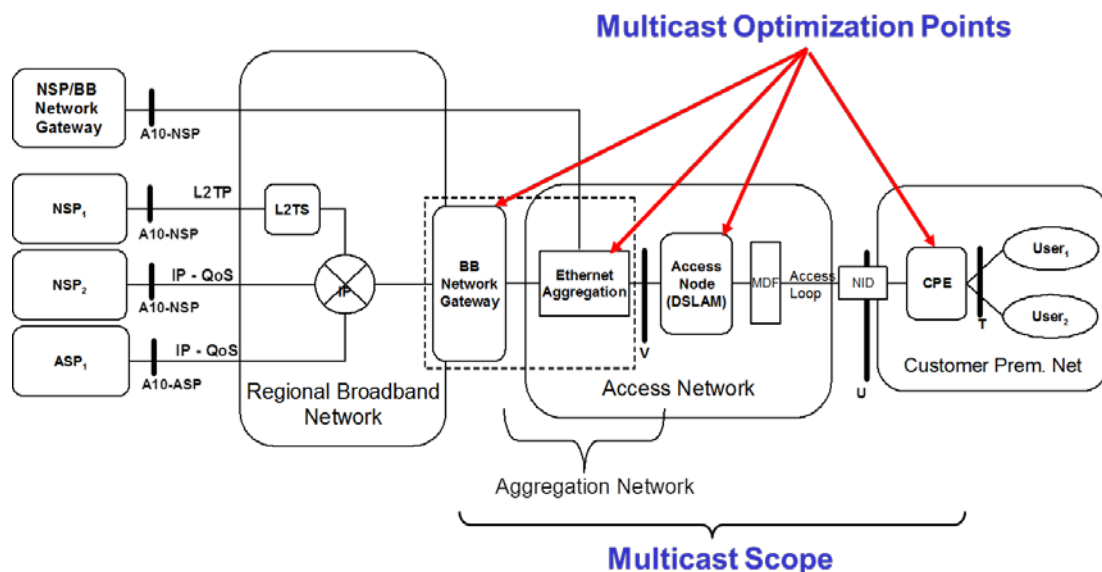


Figure 20: Multicast Reference Model from BBF TR 101

Multicast can be decomposed into three functions - multicast replication, replication control and multicast group control. The L2 WAP L2 Ethernet based network needs to support multicast optimization by controlling the flooding of Ethernet multicast frames using IGMPv3 (or MLDv2 snooping [RFC4541] for IPv6), such that packets are replicated only on those ports (physical and logical) that have specifically requested a multicast group. The L2 WAP network uses Ethernet virtualization in order to offer multicast for wholesale services. The virtualization is implemented using a L2 VLAN based approach where one or more N:1 VLANs can be used to carry IGMP (or MLD) messages and multicast traffic for wholesale multicast services such as IPTV. IGMP messages in the upstream direction (from end-users) on a certain VLAN will set up state (MAC-address filters¹⁰) in bridges to limit the flooding of multicast data in the downstream direction. These bridges include both aggregation switches and Access Nodes and is where IGMP snooping occurs. Based on the requested IP multicast group (discerned by snooping IGMP packets), the L2 WAP network will set up a L2 multicast filter entry that

¹⁰ Note that these filters are applied to reduce bandwidth. Without them a bridge will flood multicast out on all ports that are part of that multicast VLAN.



allows or prevents packets to flow to the port on which it received the IGMP or MLD report. The L2 WAP network should support multicast replication control on a per VLAN basis. The L2 WAP network's Access Nodes should support proxy reporting as per TR-101. This function actively filters IGMP packets in order to reduce load on the ANO's multicast routers. Multicast joins and leaves heading upstream to the ANO's multicast router are filtered so that only the minimal quantity of information is sent, regardless of how many active multicast listeners are connected to the Access Node.

The L2 WAP network should not inhibit use of both Any Source Multicast (ASM) and Source Specific Multicast (SSM). The L2 WAP network also needs to facilitate support for both IPv4 and IPv6 multicast. Controlling the amount of multicast traffic that can be replicated across a specific interface is a capability that the L2 WAP infrastructure provider would wish to have for traffic and capacity management purposes. In order to control jitter performance on multicast streams, ANOs can shape the streams on a per channel basis (as opposed to simply shaping the aggregate of the multicast channel streams). The ANO may also have to shape the sum of the multicast VLAN/stream and unicast VLAN traffic to ensure that it fits within the peak downstream bandwidth of the end-user's broadband access connection.

It would be desirable to have the ability to increase the default number of supported multicast groups per port per customer (generally 8) to a higher level (e.g. 16) on the L2 WAP provider's network. Dependent on the type of IPTV solution delivered, the multicast requirements may differ from the default support. For example, 4 STBs in a home all requiring a multicast feed and all STBs support Picture-In-Picture (a multicast feed per preview channel on the screen) may exceed the default supported multicast groups per customer.

The L2 WAP Network Provider's product collateral should give a clear explanation of the multicast control functionality that the ANO can exploit.

L2 WAP MULTICAST BEST PRACTISE

- ✓ **Multicast Frame replication functionality for Local and Regional handovers on dedicated N:1 VLAN**
- ✓ **IGMPv3 snooping for end-user access control of multicast in accordance with BBF TR-101 (& MLDv2 for IPv6)**

Best Practise Key Points 5: Multicast

5.5 Security & End-User Identification

The wholesale L2 WAP network connectivity product should provide sufficient end-user isolation and service integrity to prevent any individual end-user initiating a denial of service attack, or theft of service from the network or another customer. The L2 WAP network must also prevent inappropriate leaking of information or content between end-users.

Mechanisms should be inherently included to protect the network against distributed denial of service attacks, e.g. by rate limiting various types of control packets. Another is the support of Access Control (e.g. by specifying lists or filters). For L2 WAP, anti-spoofing mechanisms are primarily required at L2 but



are also required at L3 (IP/IPv6) where systems such as B2B portal web servers form part of the service offered to ANOs. Mechanisms that support both security and privacy include L2 separation of both traffic types and customers, and the prevention of basic traffic diversion techniques such as hair-pinning between customers at the Access Node [10].

End-user security is also important and has a strong overlap with privacy. Again, L2 separation of customer traffic and preventing hair-pinning are key capabilities that should be inherent in the L2 WAP network design. Using VLAN tags to uniquely identify customers on the NNI allows the separation internet in the IEEE specifications to deliver this L2 separation.

As a fundamental principle, it must not be possible that anything that one ANO or its end-users can do will affect another ANO or its end-users. That is, a malicious customer of ANO A must not be able to see or affect the traffic of any customer ANO B.

The L2 WAP network should also be transparent to traffic including any form of tunnelling at layers above¹¹ the Ethernet layer (including use of IPv6 connectivity). Encryption is one of the main ways of ensuring the privacy of customer data, and protection of commercial content (e.g. DRM). These are normally done end to end, and at the application layer. The L2 WAP infrastructure should ensure that these techniques can be used transparently across the L2 WAP network.

The L2 WAP Network Provider should also offer a means for the ANO to uniquely identify their end-users. This helps to automate provisioning and to assign service attributes (including IP address) and service entitlements. The preferred approach is to know the physical port and Access Node identifiers to which the end-user is connected since this provides a schema whereby an end-user's "co-ordinates" are unique within the L2 WAP provider's network. This can then be communicated to the ANO via DHCP¹² option 82 (or the equivalent LDRA option 18/37 for IPv6) or PPPoE Intermediate Agent techniques [4]. The Vodafone preferred E-NNI VLAN mapping will also allow customers to be uniquely identified by the VLAN tags on the NNI. This VLAN approach can be used in conjunction with DHCP Option 82 or PPPoE, or on its own where appropriate. An example of best practise is where the ANO can include their own customer ID reference in the broadband access connection order they place to the L2WAP Provider. Then, where the L2 WAP Provider does use DHCP Option 82 or PPPoE IA to communicate the circuit ID (e.g. Access Node/port) they can also include the ANO's customer ID in the message. Hence this binds the customer and circuit reference and helps facilitate a better "zero-touch" provisioning process.

¹¹ Within L2, Ethernet control protocols such as PAUSE/RESUME etc. need not be carried, as per MEF standard practice. The VULA Network Provider may look at the Ethertype and use 0x8100 or 0x88ad to determine the VLAN, DEI and Priority. It must not look deeper and reject frames based on nested 0x8100 for instance.

¹² Consideration needs to be given to any exception situations where non-DHCP IP addresses such as static addresses are used.



NOTE: In a situation where the ANO may use multiple wholesale L2 WAP Network Providers, the ANO should ensure that the customer identification schema is able to avoid the same ID being given for different end-users by the different wholesale network providers. Also, some protocols like VRRP and multicast re-use MAC addresses, implying that in some situations it needs to be possible for the end-users of one ANO to use the same MAC address of other end-users of the same (or a different) ANO, otherwise the potential for a Denial of Service (DOS) attack exists. The L2 WAP network and associated processes needs to be designed to cope with these scenarios.

SECURITY & END-USER IDENTIFICATION BEST PRACTISE

- ✓ **Customer Identification by Access Node and physical port identifiers (which may then be communicated via DHCP option 82 or PPPoE IA) - VLAN identifiers are an alternative only in situations where they provide unique customer identification co-ordinates.**
- ✓ **MAC address anti-spoofing - duplicate MAC address detection and rejection of traffic from duplicate MAC address sources¹³**
- ✓ **Control and policing of IGMP rate for N:1 VLANs**
- ✓ **Rate limit Layer 2 broadcast**

Best Practise Key Points 6: Security & End-User Identification

6. Service-Wrap Characteristics

6.1 B2B & Portal Interfaces

A B2B interface and a portal interface should be provided to facilitate transactions between the L2 WAP Network Provider and the ANO. This is required in order to support automated processes for activities such as booking appointments, problem handling, raising faults, pre-approved costs and billing. A portal approach is often used in the early days of deployment before volumes justify full B2B integration, or when new product variants are being trialled prior to launch. The systems supporting the B2B and portal automation should have adequate capacity and availability to support SLAs for responsiveness and uptime. The B2B and portal interfaces should also support an availability checker which communicates regular updates of NGA network build coverage. The B2B interface should be clearly defined and documented. It should ideally be based a flexible technology such as XML that facilitates simple, rapid and cost-effective development of enhancements. The B2B gateway needs to be secure, only enabling access by authorised users (e.g. via use of certificates etc.).

¹³ This may necessitate an agreed process to make sure that moving CPE within the network does not result in loss of service e.g. when a customer moves house and takes their CPE (and hence its MAC address) with them - Preferably without imposing a limit on the number of MAC addresses an end-user may utilise without agreeing any such limits with the ANO.



B2B AND PORTAL BEST PRACTISE

- ✓ **Automation of all key process interactions between L2 WAP Network Provider and ANO via B2B and portal interface options**
- ✓ **SLA on systems availability and response time**

Best Practise Key Points 7: B2B & Portal

6.2 Provisioning

To be able to provide service to a premise, an agreed address format for use on a web or B2B interface is needed; this is usually based on some standardised address system for each local country, with sufficient detail to identify the exact premises, including flats or apartments, businesses, and non-served premises such as cell sites or street furniture. The use of a regularly updated generalised database or flat-file can help in planning purposes for e.g. Regions, COs being enabled, roads or cabinets serving end users week by week forecast. However, accurate order management processes are needed to handle address errors, or landlord or wayleave/right of way issues that may hold up installation.

Once the correct premises is identified, then the order needs to select and confirm configuration options such as type of ONT, VLAN IDs for each customer, copper line profile to map onto initial DLM setting, etc. Processes should define sufficient messages via the B2B or web interface¹⁴ that allow both parties to understand where the order is in the 'plan and execute' phase, and allow appointments for access to be made in advance including discussion with the end-user to confirm their availability, if they are needed. There should be flexibility to allow some changes before order delivery, such as the configuration options, up to a 'Point of No Return' e.g. changing an ONT type may need to be prevented once the order is within 4 days from install, but a VLAN change could be allowed up to the day before the install.

For provision of DSL-based L2 WAP it is desirable to be able to select DSL line profiles or to set DSL line profile parameters. It is also desirable to be able to select the stability parameters in any DLM system used. These options enable the ANO to trade-off speed, stability and latency in order to optimise the line's performance to suit the requirements of their end-user. E.g. a residential game may want a profile optimised for speed and low latency where as a business user connecting an IP-PBX may prefer the line to be optimised for stability.

For provision of FTTH (e.g. GPON) L2 WAP to multi-tenancy units (e.g. apartment buildings) it is desirable to have an option for the ANO to self-provide the fibre drop within the building from the patch-panel or last splitter. This can facilitate operational process improvements in the provisioning interactions with end-users.

¹⁴ The processes associated with provision of the VULA service should be transparent with timely reporting of the progress and status at each step together with corresponding KPIs.



Initial network build will also require ordering and provision of the Interconnects. Capacity management may also identify the need for additional capacity as demand grows, and there should be clear processes including timescales to both provide the initial link, and augment with extra ports. As discussed previously, 1G and 10G options may be offered, and incremental handovers could be 1G or 10G. The network provider should clearly identify any downtime for upgrades, especially for moving from one port to multiple in a LAG group.

Migration processes should also be considered and designed as a variant of the provisioning process (including capabilities such as DSLAM or PON OLT port-change). Transferring customers between ANOs is potentially problematic from the end-user's perspective. There needs to be a robust process in place that prevents 'slamming' but which also ensures that the end-user is not without service for days or weeks owing to a lack of coordination between the various parties. For example, to avoid the case where the end-user transfers service but between their request and implementation of the re-provisioning, the resource (e.g. port on a DSLAM) has been re-allocated leaving the end-user without service.

PROVISIONING BEST PRACTISE

- ✓ **Clear processes for provision and handling of errors or changes**
- ✓ **Individual customer order progress reported regularly during all phases of provisioning**
- ✓ **All key customer provide status milestones automatically notified via B2B interface and portal**
- ✓ **Interconnect planned for growth, minimal upgrade impact on live traffic**
- ✓ **Planning information on NGA rollout/coverage (by L2 WAP Network Provider) provided on a timely and regular basis**
- ✓ **Ability to select and configure DSL line profiles/parameters and DLM stability thresholds**
- ✓ **Option to self-provide the fibre drop within multi-tenancy buildings for FTTH**

Best Practise Key Points 8: Provisioning

6.3 Assure

The L2 WAP network needs to support the ability for each ANO to use their own OAM scheme. This may require per service and per (virtual) link OAM, which work independently of the underlying transport and/or virtualization technology. The per service OAM capabilities will be used to monitor end to end continuity of the service, as well as verifying connectivity to UNI interfaces and testing the path towards the UNI interfaces. The per-service OAM capabilities can also be leveraged to carry performance management data, if necessary.

At L2 WAP hand-off points (UNI, NNI), there should be ways to provide Local Management interfaces to for example, CPE devices. In the case of a MEF UNI, these Local Management interfaces can be leveraged to provide status messages of the service availability, as well as providing information about the service itself (such as service instance tags used, traffic profiles, etc).



One of the most important areas of standardization for realizing the vision of a retail CPE model described in section 5.2 is that of Operations, Administration & Maintenance (OAM). Key issues include conformance to standards together with maturity and interoperability of vendor solutions. This also underpins assure processes and interconnect commissioning. Standardised OAM functionality is illustrated below (from [9]):

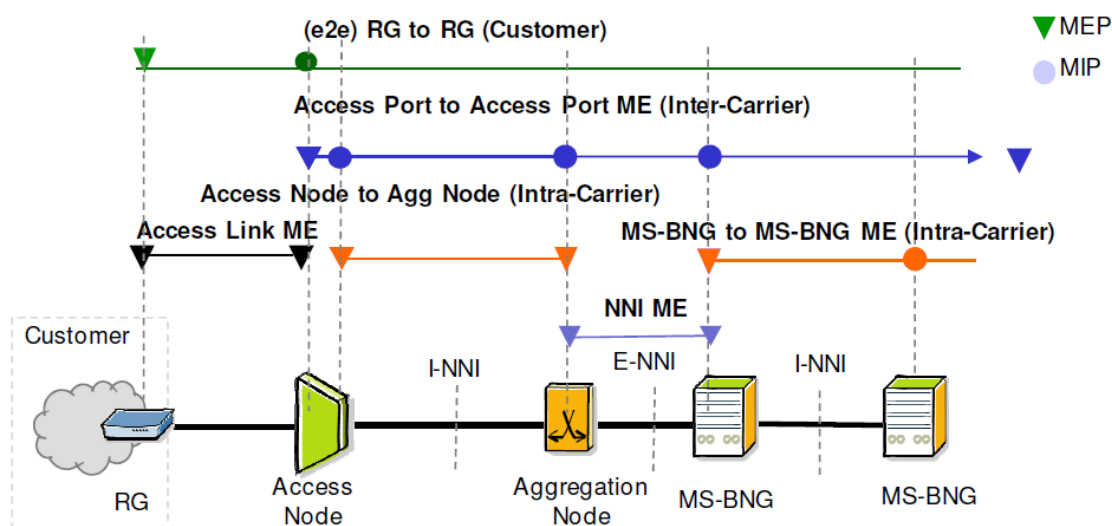


Figure 21: OAM Functionality within L2 WAP Ethernet Broadband Architecture

The “wires-only” CPE model requires fault demarcation between the customer premises and the Access Loop. Ethernet OAM between the ANO and the L2 WAP network operator across the access loop provides part of a solution for this. This requires support for Ethernet OAM functions on the Access Loop interface in both the L2 NT and the Access Node. In order to support a premium SLA (e.g. for business end-users), Ethernet performance monitoring should be supported at either the L2-NT or the Access Node.

Once diagnosis has identified the problem, B2B and web portal methods should be available to pass and progress the fault back to the network provider for repair. These electronic interfaces should also be used to convey timely and regular status updates on the progress of the repair process.

Regular reports on network health and performance levels should be provided to the ANO. Ethernet OAM supports frame-loss measurement (ETH-LM) and frame delay measurement (ETH-DM). These measurements can be used by the L2 WAP wholesale network provider and the ANO to generate values for the frame-loss, frame delay and frame delay variation described in the QoS Service Level Agreement (SLA).

In addition to Ethernet L2 OAM, the L2 WAP product should also facilitate test and diagnostics capabilities such as troubleshooting commands invoked via a SOAP/XML B2B interface or web portal. Capabilities should include a physical layer status check (e.g. DSL line in synch with DSLAM, downstream/upstream loop loss), copper line test (metallic layer tests), optical layer status (laser light levels and OTDR test result access) and remote NTE re-boot. The capabilities should include interactive



and batch test options plus the ability to access historic alarm and test results. The L2 WAP Network Provider should provide the ANO with visibility of the Access Node's port status, speed, parameters set etc. for the access connections used by the ANO. Changes to the Access Node ports assigned to the ANO should be handled real-time to allow efficient network operations and customer support processes.

ASSURE BEST PRACTISE

- ✓ **Automated capability for confirming actual configuration of provisioned parameters (line profile, VLAN configuration etc.) and performance measures**
- ✓ **Diagnostic capability for L1 and L2 testing (including loopbacks) for analysis of end user connections and for interconnects**
- ✓ **Access and management to end user CPE from ANO network via industry standards based in-band methods**
- ✓ **All key customer repair status milestones automatically notified via B2B interface and portal**

Best Practise Key Points 9: Assure

6.4 SLA/SLG

Service Level Agreements (SLAs) set out the supplier's commitment to provide a service to an agreed quality, (e.g. within x days). Service Level Guarantees (SLGs) specify the financial penalty for not meeting a particular target. Our ability to compete and to offer services to our customers depends on the quality of the input we source from the L2 WAP supplier. SLAs & SLGs are critical elements of reference offers / wholesale contracts. Robust SLAs & SLGs give us visibility and predictability of the various stages and associated timings (e.g. ordering) of wholesale fixed services.

Effective and enforceable SLAs & SLGs provide incentives to suppliers to honour their commitments - they enable us to meet our customers' requirements by setting appropriate SLAs & SLGs at the retail level. SLAs should focus on ordering, provisioning, fault-finding and restoration, and availability or up-time of the service. SLAs should apply to the electronic systems and gateways offered by suppliers, since without access to place or check progress orders, or raise and chase faults, we cannot manage our end-user service. That said, incumbent suppliers have been reluctant to include SLAs backed up by SLGs.

The following are key areas that need to be considered as potential deviations from "best practise":

SLAs:

- Incomplete SLAs: SLAs not covering all key elements of wholesale services
- SLAs but no associated penalties
- SLAs expressed as "targets" or "aims" rather than enforceable commitments
- SLAs that set too generous commitment levels / commitments which do not meet our retail customers' requirements
- SLAs that do not apply in 100% of cases



SLGs

- SLGs which have too low maximum penalties caps and/or limited escalator and/or when claimed preclude a claim for additional loss, thereby providing no or reduced incentive to rectify poor performance
- SLGs that are not high enough and do not provide incentive to comply with the SLAs

Measurements

- SLAs / SLGs that apply at some aggregate level (i.e. across all circuits) and not per line/circuits which is what matter
- Service availability poorly defined (i.e. not capturing all the facets of a functioning service)
- SLAs for service availability measured at an aggregate level and/or over a too long period (e.g. annual, whereas we are typically invoiced on a monthly basis)

Carve-outs

- Wording which includes too many carve-outs (e.g. force majeure, stop the clock) rendering the SLAs – SLGs regime ineffective

Note also that an SLA/SLG per customer access connection is not a valid solution in the case of a major network fault that impacts multiple customers/connections. Hence the per-customer SLA/SLG approach should be complemented with SLAs and penalties associated with major network and systems outages that may affect multiple connections, and also the inclusion of KPIs that reflect the performance of the overall product (network, systems & processes) such as availability and average plus 90 percentile provide and fix times etc.

Summary of Best Practise for SLAs & SLGs:

Order Validation	<ul style="list-style-type: none"> • Reasons for rejection should be provided (missing or incorrect info making it not feasible to process) • Where an order cannot be met, the incumbent should offer an alternative (e.g. for duct/pole access, the incumbent could provide an alternative routes as a matter of course) • For complex orders/services (e.g. leased lines), it may be appropriate to break down the order process into a number of stages (confirmation of receipt, site survey, etc)
Provisioning	<ul style="list-style-type: none"> • Where construction affecting provisioning time is required, different scenarios & associated timeframes should be defined in ROs. Clear rules and justification should be provided when longer timeframe / construction costs apply. Evidence of expected construction cost should be provided • Compliance with the SLA for provisioning should be measured against the incumbent's agreed delivery appointment date, not against delivery within the maximum Provisioning timeframe. • Situations where 'Stop the Clock' is permitted should be clearly defined • Where the applicability of SLAs/SLGs for provisioning are conditional on meeting forecasting requirements, those should not be overly restrictive (e.g. SLAs/SLGs may not apply for circuits in excess of 30% of the forecast)
Service availability & QoS	<ul style="list-style-type: none"> • Measurement period should be monthly or quarterly at most • Detailed performance reports should be provided by incumbents on a monthly/quarterly basis • Separate SLAs with detailed parameters for service availability & QoS (e.g. jitter, packet loss for leased lines) should be defined
Fault restoration & degradation of quality	<ul style="list-style-type: none"> • There should be clarity over whether a fault SLA relates to complete outages only, or also to service degradation. If the former, a separate degradation/QoS SLA should also be set. • An SLA regarding the frequency with which faults can be raised against a specific circuit (to prevent tickets on a broken circuit being incorrectly repeatedly closed to 'game' the Fault



	SLA) should be considered. <ul style="list-style-type: none"> The SLA should be measured as a 100% commitment for each individual fault, not measured on an average success rate.
--	--

Table 3: SLA & SLG Requirements

For L2 WAP, we require SLAs for provision and repair of the NNI handover interconnect port(s), as well as the end-user connections, and for service-flow related features such as adding additional QoS flows, multicast etc. Ideally there would be SLA on availability – 99.9% etc on individual services, with packet loss measures. These need to reflect the tightest service parameters that are envisaged. For example, mobile backhaul via L2 WAPs may have to be a trade-off between the QoS Services offered vs the timing parameters of the different mobile services – the table below shows the parameters of interest:

Service	Technology	Downstream Mbps	Upstream Mbps	Latency one way (e2e) ms	Latency oneway (vdsi) ms	Packet Loss	Jitter (ms)	Frequency Sync req. (ppb)	Phase Sync req. (us)	Time Sync (us)	Sync Protocol
2g		20/50	10 or more	40/60		<0.05%	15/20	+/- 50	Not required	Not required	1588v2 / Synch E preferred
3g	UMTS	20/50	10 or more	40/60		<0.05%	15/20	+/- 50	Not required	Not required	1588v2 / Synch E preferred
4g	LTE FDD	100 -> 400	30 or more	Recommended 10 Tolerable 20	Huawei lab test data, FYI: DS Latency : 0-3300 us, AVG DS Latency: 1200 us, US Latency 0-5700 us, AVG US Latency :1300 us [0-400 meters, 4*VDSL2 bonding with mix length traffic, 100% load, no noise]	Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	Not required	Not required	1588v2 / Synch E preferred
	eMBMS	100 -> 400	30 or more	Recommended 10 Tolerable 20		Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	+/- 1	+/- 1	1588v2 / Synch E preferred
	Network MIMO	100 -> 400	30 or more	Recommended 10 Tolerable 20		Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	+/- 0.5	+/- 0.5	1588v2 / Synch E preferred
	ICIC	100 -> 400	30 or more	Recommended 10 Tolerable 20		Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	+/- 1	+/- 1	1588v2 / Synch E preferred
	COMP	100 -> 400	30 or more	Recommended 10 Tolerable 20		Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	+/- 0.5	+/- 0.5	1588v2 / Synch E preferred
	Location Service	100 -> 400	30 or more	Recommended 10 Tolerable 20		Recommended <0.001% Tolerable<0.5%	Recommended 4 Tolerable 8	+/- 50	+/- 0.2	+/- 0.2	1588v2 / Synch E preferred
Lampsite	LTE + UMTS	Minimum: 21Mbps(1 UMTS cell) Medium: 200Mbps (2 UMTS + 1 LTE) Long term: 350Mbps (2 UMTS + 2LTE cells)	Minimum: 5 Mbps(1 UMTS cell) Medium: 60Mbps (2 UMTS + 1 LTE) Max: 40ms (2 UMTS + 2LTE cells)	Target: 10ms		Max: 0.001 Target: 0.0001	Max: 15ms Target: 2ms	+/- 50	Not required (+/- 1.5 for LTE)	+/- 1.5	1588v2 / Synch E preferred

Table 4: Mobile Backhaul Requirements Pertinent to NGA

SLA/SLG Best practice

- ✓ SLAs & SLGs should be defined for each of the main elements of the life cycle of services including at least, Ordering, Provisioning, Service Availability & Fault restoration
- ✓ SLAs & SLGs should be defined for the electronic platform used to interface with the L2 WAP Network Provider, including availability and response times.
- ✓ SLAs should be aligned with end-user requirements
- ✓ SLGs should be high enough to incentivise compliance with the SLAs by the L2 WAP Network Provider, preferably with no penalty caps and with a right to claim for additional losses above the level of SLGs
- ✓ SLAs & SLGs should apply per fault/event/line/circuit - not in aggregate for average performance
- ✓ Payment of penalties should be pro-active /automatic – The ANO shouldn't have to measure it or ask for it



- ✓ The L2 WAP Network Provider should provide reports on actual performance against SLAs – ANOs should have a right to challenge reported performance with contrary evidence
- ✓ NRAs should collate and publish incumbent service performance
- ✓ SLAs & SLGs should be tightly worded with limited carve-out conditions, clearly identify exceptions, limited opportunities for stop-the-clock, outage time for electronic platforms and other multiple feed-back loops

Best Practise Key Points 10: SLA/SLG

7. Evolution Considerations

7.1 Access Transmission System Evolution

The L2 WAP architecture and capabilities described in this document are focussed around the end-to-end Ethernet L2 connectivity. As such, this is relatively independent from underlying transmission system used to provide broadband access. Hence, as broadband access technologies evolve, the L2 WAP architecture can be evolved to incorporate such technology developments. L2 WAP network capabilities have been deployed for VDSL and GPON access technologies for a number of years. These can be evolved quite easily to accommodate enhanced and vectored VDSL, bonded VDSL, G.fast and NG-PON2 access technologies.

Faster fixed access technologies increase the potential range of L2 WAP uses beyond broadband access for consumers and business end-users. Mobile backhaul over broadband access will become increasingly important as small-cells proliferate for 4G and 5G mobile networks. In addition to the prerequisite bandwidth, the L2 WAP network also needs to be able to support synchronisation capabilities (time, frequency and phase).

As the end-user broadband speeds increase with the faster access technologies it is anticipated that L2 WAP NNI interfaces will increasingly need to evolve from 1 Gbit/s to 10 Gbit/s. The delivery of Ultra-HD TV (4k) over these NGA technologies could also increase the focus on use of the multicast capability of L2 WAP. L2 WAP products should enable upgrades (e.g. from VDSL to G.fast or GPON) as the footprint of newer, faster NGA technologies expands and potentially overlaps with “legacy” technologies. Migration options to newer NGA technologies together with associated processes and B2B automation should be supported as part of the L2 WAP product capabilities. In addition, for general hardware and software evolution upgrades (that don’t necessarily mean a technology change) it is desirable for the L2 WAP Network Provider to have a test environment for validation activities such as checking that CPE interoperability continues with the new “network release”. This will help minimise service outages.

7.2 The Impact of Virtualisation

Wholesale broadband access has existed for over 15 years now in terms of both business models and network deployments. In the early days of broadband, the network infrastructure providers (in many markets) sold wholesale broadband access connectivity products to numerous retail ISPs. Many of these wholesale networks “virtualised” their network and OSS infrastructure by allowing the ISPs a degree of access to their own end-user’s broadband lines. Initially this was focussed on Test & Assure capabilities. This enables the ISPs to provide more effective first line support and effectively outsourced some of the troubleshooting from the wholesale network provider to the ISP. The ISP could



access “a slice” of the wholesale network provider’s systems to test and diagnose DSL lines. Examples of the sort of capabilities provided include [22], [23]:

- User session status (e.g. PPP session information) and RADIUS log check
- IP ping tests & packet counters
- ATM layer 2 F4/F5 loopback tests and cell counters
- ADSL status check (line in synch with DSLAM, downstream/upstream loop loss)
- Copper line test (metallic layer tests & probable root cause of faults)
- Interactive and batch test options
- Historic alarm and test result access

Secure access to the appropriate wholesale systems was usually provided via a secure authenticated gateway portal.

This “virtualisation” of the wholesale provider’s network and systems then evolved to include configuration capabilities such as the ability to select the DSL Dynamic Line Management (DLM) line stability option (e.g. fast, stable, super-stable ..) or even explicitly allowing the ISP to select individual DSL profiles (see [23]). Hence virtualisation has been a part of wholesale broadband for many years now and overtime has evolved from offering ANOs (such as ISPs) just test and diagnostics functionality to now including a degree of configuration capability. This trend looks set to continue and it is an obvious way in which L2 WAP products could be improved to allow the ANO’s a greater control over features with which to develop differentiated retail NGA services. Consequently this will result in a greater choice for end-users.

In summary, it has been proven in existing deployments that it is entirely feasible for wholesale providers to provide access to a constrained “slice” of their OSS and element management systems to safely share data such as from an Access Node. For example providing federated access to a 3rd party (the ANO) to DSLAM data and constrained access to the wholesale providers control plane, either for monitoring, or control (or both). Such slicing of the NGA network management systems can be applied to existing NGA network deployments and so does not necessitate the upgrade or change-out of network equipment elements (e.g. to include virtual machine hypervisors etc.). Hence the approach does not require radically new and complex systems development that would be significantly beyond what has been historically provided. The systems implications, business models, roles, responsibilities, commercial relationships and obligations are simply variants of those addressed by the original wholesale models [22], [23]. A new generation of NGA equipment that does include enhanced virtualisation capabilities (such as virtual machine segmentation within the hardware) may bring further capabilities that lead to improved L2 WAP products offering greater control for ANOs. However, this is not essential for an initial “Network as a Service” approach since management systems segmentation offers a valid first step.



8. Prospects for L2 WAP Product Alignment Across Markets

Given that L2 WAP products with differing characteristics have already been launched in a number of countries [2], it is worth considering the practicalities of retrospectively altering them to create a more aligned approach. This section considers each key L2 WAP product characteristic in turn from the prospective of changing them to create a more internationally harmonised L2 WAP product:

- Network Architecture & VLAN configuration

This is potentially the most difficult in that it is not desirable to change the VLAN configuration or numbering at the UNI in a way that impacts the interoperability of existing deployed customers so that they go “off air”. Altering the architecture by addition of new VLAN connection options may be more viable so that existing CPE continues to function, but just can’t take advantage of new options. Changes to connectivity within the network can to an extent be hidden from the ANO as long as the UNI and NNI are not impacted. Changing an existing deployment from an N:1 VLAN model to a 1:1 VLAN model is more substantial and impactful.

- Network interconnect location options

If the L2 format of the NNI is preserved then it should be feasible to add (or remove) backhaul connectivity options to increase the range of NNI locations offered and include migration options too.

- User equipment options – CPE/modem choice

Most markets that offer a wires-only UNI enabling the ANO to provide their own CPE started with the CPE being provided by the L2 WAP Network Provider. The Broadband Forum interoperability standards and certification programmes greatly facilitate the ability to move to a wires-only model for both DSL and fibre L2 WAP products.

- Bandwidth, traffic prioritisations and QoS

The ability to prioritise bandwidth and provide QoS SLAs is mainly determined by the hardware capabilities of the Access Nodes and any associated aggregation devices. As long as there are not overly onerous limitations on the hardware such as limited queues or scheduler capabilities then adding appropriate assured bandwidth, prioritisation and QoS is largely a matter of configuration (after extensive testing) and enhancing the associated capacity management processes.

- Multicast

Multicast is another capability where the challenge in retrospectively adding it depends largely on the capabilities of the hardware elements already in situ. Nearly all modern Access nodes and Ethernet aggregation switches can support multicast and a number of L2 WAP Network Providers have introduced such a capability subsequent to initial L2 WAP product launch.

- Security & End-User Identification

Network equipment that is compliant with the requirements of the relevant broadband architecture standards ([1], [4], [6]) should already be capable of supporting the suggested best practises. Hence it is mainly a question of configuration and testing in order to deploy these.



- B2B & portal interfaces

It is difficult to conceive of a viable L2 WAP product that would not use electronic interfaces (as front-ends to largely automated processes) in order to scale effectively. B2B interfaces and schema are normally already on a regular release upgrade cycle to accommodate new products and other such enhancements. The ability to align approaches will depend to an extent on the underlying technologies deployed. Modern interfaces based on Internet approaches such as XML and SOAP etc. will be easier to change.

- Provide process

Any L2 WAP Network Provider with a culture of customer focus and continuous improvement should already be able to adjust and optimise their processes for provision of service. Some systems development for improved automation may be required within some L2 WAP Network Providers in order to bring this capability up to the levels of "best practise".

- Assure process

As above, any L2 WAP Network Provider with a culture of customer focus and continuous improvement should already be able to adjust and optimise their processes for service assurance. Some systems development and deployment of additional technologies such as probes may be required here.

- SLA/SLG

This can start very simply by measuring existing process and network performance bounds and variation and adding an appropriate a service wrap. However, this would likely result in a mismatch with user requirements. Hence additional measurement and reporting processes and systems development are likely to be required. The most significant aspect of retrofitting best practise in this area could be commercial in terms of creating and validating new contract terms and conditions etc.



9. Abbreviations

ADSL	Asymmetric DSL
AF	Assured Forwarding
AN	Access Node
ASP	Application Service Provider
B2B	Business to Business
BBF	Broadband Forum
BE	Best Effort
BNG	Broadband Network Gateway
CIR	Committed Information Rate
CO	Central Office
CPE	Customer Premises Equipment
DEI	Drop Eligibility Indicator
DHCP	Dynamic Host Configuration Protocol
DLM	Dynamic Line Management
DOS	Denial of Service
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexor (an Access Node)
EF	Expedited Forwarding
E-NNI	External NNI
FANS	Fixed Access Network Sharing
FTTC	Fibre To The Cabinet
FTTdp	Fibre To The Distribution Point
FTTH	Fibre To The Home
GPON	Gigabit Passive Optical Network
HGU	Home Gateway Unit (i.e. GPON ONT variant)
IGMP	Internet Group Multicast Protocol
IVC	Infrastructure Virtual Circuit
L2TP	Layer 2 Tunnelling Protocol
LDRA	Lightweight DHCPv6 Relay Agent [RFC 6221]
LLU	Local Loop Unbundling



MEF	Metro Ethernet Forum
MLD	Multicast Listener Discovery
MSAN	Multi-Service Access Node
MS-BNG	Multi-Service Broadband Network Gateway
NGA	Next Generation Access
NID	Network Interface Device
NNI	Network to Network Interconnect
NSP	Network Service Provider
NT	Network Termination
NTE	Network Terminating Equipment
ODN	Optical Distribution Network
OLT	Optical Line Termination (GPON Access Node)
OMCI	Optical Management & Configuration Interface
ONU	Optical Network Unit (a GPON NTE variant)
ONT	Optical Network Termination (a GPON NTE variant)
PIA	Passive Infrastructure Access
PPPoE IA	Point to Point Protocol over Ethernet Intermediate Agent
RG	Residential Gateway
SBU	Single Business Unit (i.e. GPON ONT variant)
SLA	Service Level Agreement
SLG	Service Level Guarantee
SPR	Snooping with Proxy Reporting
STB	Set Top Box
UNI	User to Network Interconnect
VDSL	Very High Speed DSL
VID	VLAN Identifier
VLAN	Virtual Local Area Network
VRRP	Virtual Router Redundancy Protocol
L2 WAP	Virtual Unbundled Local Access
WAN	Wide Area Network



10. References

- [1] Broadband Forum TR-178, "Multi-service Broadband Network Architecture and Nodal Requirements", Issue 1, September 2014.
- [2] BEREC, "Common Characteristics of Layer 2 Wholesale Access Products in the European Union", June 2015.
- [3] NICC ND1030v1.1.1, "Ethernet ALA Service Definition", 2010.
- [4] Broadband Forum TR-101, "Migration to Ethernet Based DSL Aggregation", Issue 2, July 2011.
- [5] The Journal of The Institute of Telecommunications Professionals, Volume 1 Part 3 (page 33), "Requirements for Wholesale GPON (Gbit Passive Optical Network) Access", Gavin Young, 2008.
- [6] Broadband Forum TR-156, "Using GPON Access in the context of TR-101", Issue 3, November 2012.
- [7] Broadband Forum TR-144, "Broadband Multi-Service Architecture & Framework Requirements", Issue 1, August 2007.
- [8] NICC ND1642v1.1.1 "Requirements for Ethernet Interconnect and Ethernet ALA", 2010.
- [9] Broadband Forum TR-145, "Multi-service Broadband Network Functional Modules and Architecture", Issue 1, November 2012.
- [10] NICC ND1644v1.1.1 "Ethernet ALA Architecture" 2010.
- [11] NICC ND1036v1.1.1 "Ethernet ALA NNI", 2011.
- [12] Metro Ethernet Forum MEF26.1, "External Network Network Interface (ENNI) – Phase 2", January 2012.
- [13] "Code of Conduct on Energy Consumption of Broadband Equipment" Version 2, 17th July 2007.
- [14] Broadband Forum TR-142, "Framework for use of TR-069 with PON Access".
- [15] Broadband Forum TR-069, "CPE WAN Management Protocol".
- [16] Broadband Forum TR-104, "DSL Home Provisioning Parameters for VoIP CPE", September 2005.
- [17] Broadband Forum TR-135, "Residential Data Model for a TR-069 Enabled Set Top Box", December 2007.
- [18] NICC ND1031v1.1.1, "Ethernet ALA UNI", 2010.
- [19] Broadband Forum TR-114, "VDSL2 Performance Test Plan", Issue 1, November 2009.
- [20] Broadband Forum TR-115, "VDSL2 Functionality Test Plan", Issue 1, November 2009.
- [21] Broadband Forum TR-247, "Abstract Test Plan for GPON ONU Conformance", Issue 1, November 2011.
- [22] BT Exact Technologies, "Woosh Testing for ISPs", 2002.
- [23] http://support.aa.net.uk/TalkTalk_Wholesale_Line_Profiles