TECHNICAL REPORT

TR-145

Multi-service Broadband Network Functional Modules and Architecture

Issue: 1
Issue Date: November 2012

Notice

Issue History

Issue Number	Approval Date	Publication Date	Issue Editors	Changes
1	26 November 2012	28 January 2013	A Cui, AT&T Y Hertoghs, Cisco	Original

Comments or questions about this Broadband Forum Technical Report should be directed to info@broadband-forum.org.

Editors	Anna Cui Yves Hertoghs	AT&T Cisco
E2EA WG Chairs	David Allan David Thorne	Ericsson BT
Vice Chair	Sven Ooghe	Alcatel-Lucent
Chief Editor	Michael Hanrahan	Huawei Technologies

TABLE OF CONTENTS

List of Figures

List of Tables

Executive Summary

TR-145 provides reference architectures for multi-service broadband networks, defines high level network requirements, and specifies functional modules to meet those network requirements. This Technical Report also provides high level guidance on how to distribute these functional modules onto multi-service broadband network elements. Detailed broadband network nodal requirements will be specified in further working texts e.g. WT-178.

The multi-service broadband networks/architectures specified in TR-145 support the emerging and the legacy services listed in TR-144[4]. The functional module decomposition adopted in TR-145 enables a variety of deployment options and different distributions of functions.

Convergence is a key objective driving both the content and the structure of TR-145. TR-145 specifies a converged broadband network architecture supporting many types of services, access and transport technologies, and deployment scenarios envisioned by service providers.

1. Purpose and Scope

1.1 Purpose

TR-058 [1] specified a set of mass market focused business requirements for a Multi-Service DSL architecture and evolution from the then deployed DSL architectures.

TR-144 [4] addressed a larger set of business drivers than was covered by TR-058 [1] in order to support broader market segments.

TR-145 specifies the architecture and network functional requirements to support the TR-144 [4] business requirements.

1.2 Scope

TR-145 extends the TR-101[3] architecture with new technical requirements needed to fulfill the business requirements laid out in TR-144[4] *Broadband Multi-Service Architecture and Framework Requirements*. It defines an architectural framework based upon functional modules and the logical interfaces between them, the high-level common network service requirements as well as end-to-end operational functions, such as control and OAM. It focuses on the data plane between the T and A10 reference points and provides the interfaces and connectivity to Policy Control and Management Systems via interfaces across the R and M reference points, respectively, as shown in Figure 1.

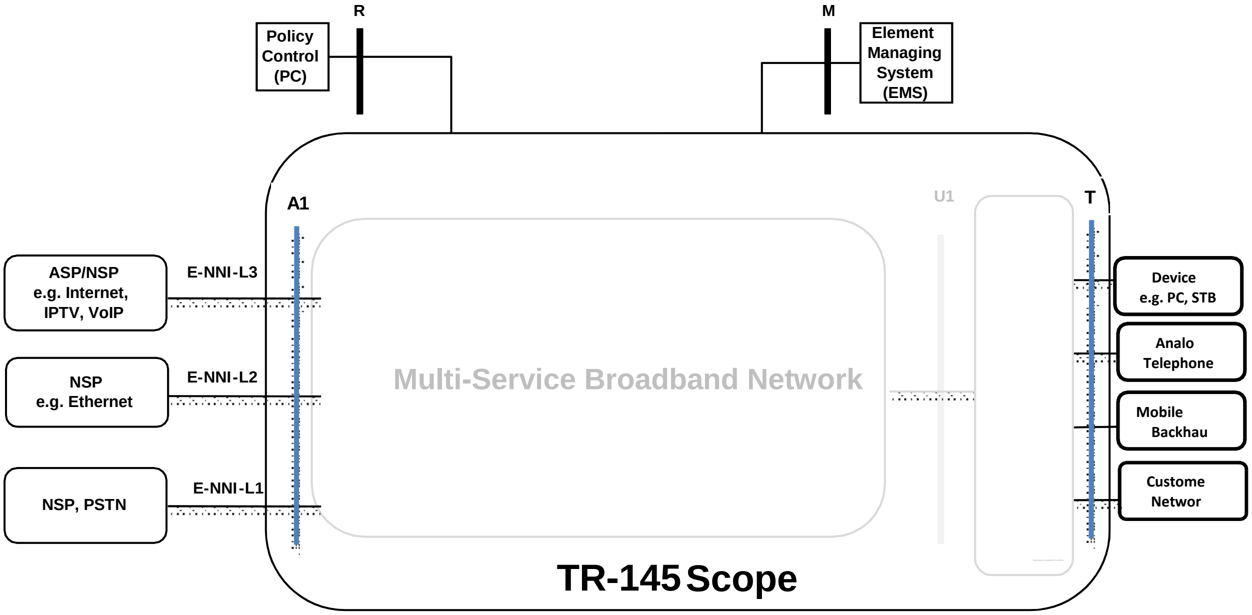

Figure 1 TR-145 Scope

There are two types of high-level functional modules defined in TR-145: technology independent, functional modules and technology specific functional modules. The latter are various instantiations of the technology independent modules. The choice of specifying functional modules in TR-145 aims at capturing as many commonalities as possible.

Any functional module can potentially be implemented on any node. Each access or aggregation node may implement only some of these functional modules or all of them, depending on the actual business requirements in specific operational contexts and timeframes. The functional modules defined in TR-145 will make it possible to have several service insertion placement strategies, on a per-service basis.

TR-145 does not redefine the market requirements laid out in TR-144 [4], but rather focuses on a functional architecture and the high-level technical requirements to support them. It will also not provide detailed nodal requirements. These will be documented in follow on working texts, in the first instance, WT-178 [14].

1.3 Relation to other Broadband Forum documents

TR-145 extends the TR-101 [3] architecture with the new technical requirements needed to fulfill the business requirements laid out in TR-144 (*Broadband Multi-Service Architecture and Framework Requirements* [4]), TR-147 (*Layer 2 Control Mechanism for Broadband Multi-Service Architecture* [5]), TR-156 (*Using GPON Access in the context of TR-101* [6]), TR-200 (*Using EPON Access in the context of TR-101* [10]), TR-167 (*GPON-fed TR-101 Ethernet Access Node* [7]), and Working Text 146 (*IP Sessions* [13]), are also considered as an existing ecosystem which TR-145 complements.

TR-145 is related to the following WTs/TRs: WT-178 (*Multi-service Broadband Network Architecture and Nodal Requirements* [14]), TR-134 (*Policy Control Framework* [12]), TR-203 (*Interworking between Next Generation Fixed and 3GPP Wireless Access* [15]) and TR-221 (*Technical Specifications for MPLS in Mobile Backhaul Networks* [11]).

TR-145 will be IPv4 and IPv6 compliant, but will leave to TR-177 (*IPv6 in the context of TR-101* [8]), TR-187 (IPv6 for PPP Broadband Access [9]), and TR-242 (*IPv6 Transition Mechanisms for Broadband Networks* [16]) the task of updating existing BBF documents to add IPv6 support.

1.4 Motivations

The motivations leading to the architecture specified in this Technical Report remain those detailed in TR-144 [4]. This section summarizes these main drivers.

- <u>Simplification of network architecture</u>: for new deployments free from migration constraints, an end-to-end architecture that is a combination of an IP/MPLS and Ethernet.

- Seamless connectivity: Integration of access, aggregation and core networks within a given administrative domain.
- A converged broadband multi-service network to support emerging as well as existing services: The broadband Multi-Service architecture must support packet based triple play services, such as residential internet access, VoIP and IPTV, as well as emerging packet based business services and packet based Mobile Backhaul. Voice services currently offered over the PSTN, TDM services used for business customers, cellular backhauling and other business applications over legacy access technologies have to be supported as well, but mainly through emulation or packetization at customer locations. To transport ATM/TDM flows over a packet network, ATM/TDM emulation over Pseudo-Wires will be utilized in the All of the above services must be supported in retail and wholesale manners concurrently.
- Operational enhancements: migration to a converged packet based access and aggregation network to minimize the number of operational touch points and to simplify and improve the end-to-end provisioning process.
- Service independence: support for separation of the client service and transport network.
- Multipoint services (e.g. for business connectivity) the network must provide the appropriate packet layer functionality in both control and data planes, typically based on Ethernet or IP/MPLS.
- Nomadism, – this allows the users to access services at different locations.
- Fixed-Mobile Convergence (FMC), which allows a user to use a single handheld device for both fixed and mobile access
- Enhanced availability with resiliency: use of OAM and an appropriate control plane for automatic convergence protection and/or restoration in case of a network link or a network node failure or maintenance activity.
- Enhanced Scalability: each IP Edge Node has a finite ARP table, and each Ethernet switch also has limited MAC table size. Distributing IP Edge Nodes closer to customers allows the connection of more customers through an aggregation network.
- Enhanced Security: avoid the risks of denial of service and other security threats acknowledged in TR-101 [3].
- Multi-Edge : the ability to source traffic to an individual subscriber from multiple Service Edges

1.5 Services supported by the network

TR-145 will not focus on application specific requirements, but the network must support at least the following applications defined in TR-144 [4]:

1 - Residential Subscriber Services:

- Data services - including all services, applications and appliances that make use of Internet access for a variety of client-server or peer-to-peer applications.
- VoIP services
- Video Services - including multicast channels and the unicast VOD service.

Support of legacy POTS

2 - Business Services

- Ethernet VPNs: E-LINE, E-TREE, E-LAN services [18][19].
- L3 VPNs
- Support of emerging converged enterprise connectivity (Voice, Data and applications)
- Support of legacy T1/E1 services (voice, data, mixed, Primary Rate Interface (PRI))

3 - Mobile Backhauling

- 2G backhauling using TDM Emulation
- 3G backhauling using ATM emulation and/or Ethernet service
- 4G backhauling using Ethernet and/or IP service
- FEMTO cell / pico-cells backhauling using any types of broadband access
- Fixed wireless IEEE 802.16 / WiMAX backhauling

4 – Wholesale and retail

The architecture needs to support all of the above services in a retail and a wholesale manner in order to comply with requirements and business needs. Wholesale and retail services are differentiated largely by the ownership of the elements in the Service Provider Domain. There is a need to allow the network operator to engage with more types of wholesale business partners. It is also foreseen that Subscriber management can be outsourced to the network provider, covering application connectivity to multiple 3rd party application providers. This 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

Note that some applications from 1, 2, and 4 above must be deliverable over a single access line simultaneously.

1.6 Access technologies

This Technical Report will not focus on access technology specific requirements, but will assume the following access technologies are supported, as listed in Annex C/ TR-144 [4]:

- xDSLDSL with and without bonding
- Ethernet
- xPON with various fiber termination locations, such as FTTP, FTTC, etc.
- Microwave wireless

- legacy voice
- wavelength access for business customers
- TDM circuits

1.7 Out of Scope

The following are out of scope of TR-145:

- Mobility specific to air interface processes, such as handovers, etc.
- Requirements, design and technical definitions for non-packet networks, such as TDM transport networks and Optical transport networks.
- All mechanisms mapping functional modules to particular nodes or addressing their coexistence in the same node are left to WT-178

2. References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [28].

MUST This word, or the terms “REQUIRED” or “SHALL”, mean that the definition is an absolute requirement of the specification.

MUST NOT This phrase, or the phrase “SHALL NOT”, means that the definition is an absolute prohibition of the specification.

SHOULD This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.

SHOULD NOT This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.

MAY This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

2.2 References

The following references are of relevance to this Technical Report. At the time of the publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below
A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

Document	Title	Source	Year
[1] TR-58	*Multi-Service Architecture & Framework Requirements*	BBF	2003
[2] TR-92	*Broadband Remote Access Server (BRAS) Requirements Document*	BBF	2004

[3]	TR-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[4]	TR-144	*Broadband Multi-Service Architecture & Framework Requirements*	BBF	2007
[5]	TR-147	*Layer 2 Control Mechanism For Broadband Multi-Service Architectures*	BBF	2008
[6]	TR-156 Issue 2	*Using GPON Access in the context of TR-101*	BBF	2010
[7]	TR-167 Issue 2	*GPON-fed TR-101 Ethernet Access Node*	BBF	2010
[8]	TR-177	*IPv6 in the context of TR-101*	BBF	2010
[9]	TR-187	*IPv6 for PPP Broadband Access*	BBF	2010
[10]	TR-200	*Using EPON in the Context of TR-101*	BBF	2011
[11]	TR-221	*Technical Specifications for MPLS in Mobile Backhaul Networks*	BBF	2011
[12]	TR-134	*Broadband Policy Control Framework (BPCF)*	BBF	2012
[13]	WT-146	*IP Sessions*	BBF	*WIP*
[14]	WT-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	*WIP*
[15]	TR-203	*Interworking between Next Generation Fixed and 3GPP Wireless Access*	BBF	2012
[16]	TR-242	*IPv6 Transition Mechanisms for Broadband Networks*	BBF	2012
[17]	MEF 4	*Metro Ethernet Network Architecture Framework - Part 1: Generic Framework*	MEF	2004
[18]	MEF 6.1	*Ethernet Service Definitions - Phase 2*	MEF	2008
[19]	MEF 10.2	*Ethernet Service Attributes Phase 2*	MEF	2009
[20]	MEF10.2.1	*Performance Attributes Amendment to MEF 10.2*	MEF	2011
[21]	MEF13	*User Network Interface (UNI) Type 1 Implementation Agreement*	MEF	2005
[22]	MEF20	*User Network Interface (UNI) Type 2 Implementation Agreement*	MEF	2008
[23]	MEF26.1	*External Network Network Interface (ENNI)–Phase 2*	MEF	2012
[24]	MEF28	*External Network Network Interface (ENNI) Support for UNI Tunnel Access and Virtual UNI*	MEF	2010

[25]	G.8260	*Definitions and terminology for synchronization in packet networks*	ITU-T	2010
[26]	G.810	*Definitions and terminology for synchronization networks*	ITU-T	1996
[27]	1588	*IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*	IEEE	2008
[28]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[29]	RFC5905	*Network Time Protocol Version 4: Protocol and Algorithms Specifications*	IETF	2010

2.3 Definitions

The following terminology is used throughout this Technical Report.

Access function set
An access function set terminates physical access media (e.g. DSL, PON, GE) provides Adaptation to the aggregation technology and forwards the traffic to a switching and aggregation instance which is part of an Aggregation Module.

Aggregation function set
An aggregation function set provides a multiplexing and switching instance which includes functions of traffic control, e.g. QoS, policy control, multicast.

Availability
Network availability is usually described in terms of 'unavailability events' which have three dimensions:

- Event frequency: how often the event occurs
- Event duration: how long it takes to restore normal operation either in terms of the average or more usefully in terms of a distribution or percentile threshold
- Event scope: how much impact a single event has (i.e. how many customers are affected at one time or the size of the geographic region affected)

The combination of event frequency, duration and scope information can be used to derive availability percentages and other more service-specific figures of merit such as lost call minutes etc.

Business Interface
An Interface between the Regional Access Provider and a 3rd party (e.g. Mobile Carrier, Business Service Customer, NSP/ASP, etc) and serves as a hand-off interface

Edge BNG A device similar to a TR-101 BNG that is part of a nested BNG architecture (see WT-178 [14])

Ethernet Service Edge A device that hosts Ethernet Services, which may or may not be session based.

Ethernet Session Analogeous to Layer 2 Session

Functional module A set of functions, which can be instantiated in a network node.

A network node can contain one or more functional modules. A functional module cannot be split between network nodes. Nodal distribution of functional modules is left to WT-178.

IP Flow As defined in WT-146, an IP Flow is defined by a 5-tuple IP parameter traffic classifier. An IP Flow can form the classification element of a traffic policy that is applied to a Session.

IP Service Edge A Service Edge that hosts IP and/or PPP Sessions.

IP Session As defined in WT-146, an IP session is a grouping of traffic according to 1 or more classifiers visible at a control point, called the IP Edge. The classifier is composed of, at a minimum a Subscriber's IP address, IPv4 subnet or IPv6 prefix (es).session for which the data plane classifier is composed of at a minimum a subscriber's IP source, including wildcards. These classifiers may be augmented by additional Layer1, Layer 2 parameters when appropriate.

L3 Session IP Session or PPP Session

Layer 2 Session Used as Part of Ethernet Wholesale Services, also Layer 2 ALA (Active Line Access) Session

Logical Interface A logical interface in the broadband architecture, at a boundary between 2 functional modules. It is shown as a line between two functional modules and a combination of letters.

Not all logical interfaces need to be instantiated as a physical interface. Several functional modules may be grouped into a single physical nodal implementation. In such a case, the logical interfaces internal to the node will not be externally visible. Hence the corresponding logical interfaces are not instantiated as physical interfaces, for example, Vc, SI-NNI.

Network node A physical, self contained element of a broadband network.

Examples: a DSLAM, an aggregation switch, etc.

Physical interface A physical instantiation of a Logical Interface. It is externally visible and may host one or several logical interfaces. Example E-NNI, V.

Reference Point A reference point is a ‘place’ inside an architecture, where one or more logical, physical, or business interfaces can be instantiated. A reference point can be internal or can be located at a given physical interface The reference points used in this document include A10, Va, Vc, U, U1, T.

Service BNG A device similar to a TR-101 BNG that terminates attachment circuits extended from the Edge BNG. (see WT-178 [14])

Service Edge A device capable of hosting subscriber sessions

Session There are four types of session:

- Access Session: this is where the access link comes up and is available for data transmission. In the DSL case, this start when the DSL modem has trained up with the DSLAM, and with ANCP the DSLAM would then transmit a Port Up message to BNG
- Subscriber Session: Layer 2 ALA Sessions, PPP Sessions and IP Sessions as defined in WT-146[24].
- Traffic rule session: this type is an abstraction of a set of policy rules. This would be used with an identifier to allow an operator to know if a particular “set of Traffic rules” is enabled or not without needing to know the underlying rule details. For example, an http redirect service would be a set of rules that allowed DNS traffic to be transmitted, redirected http traffic to a web portal, and dropped all other traffic.
- Application Session: for example a voice call, a VOD session, a gaming session or a P2P session.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AN	Access Node
A-NNI	Adaptation NNI
BNG	Broadband Network Gateway
CFM	Connectivity Fault Management
CPE	Customer Premises Equipment
DDOS	Distributed Denial Of Service
DHCP	Dynamic Host Configuration Protocol
DHD	Dual Homed Device
DHN	Dual Homed Network

EFP	Ethernet Flow Point
FCAPS	Fault, Configuration, Accounting, Performance, and Security
FSOL	First Sign Of Life
GFP	Generic Framing Encapsulation
I-NNI	Internal NNI
IVC	Infrastructure Virtual Circuit
L1A	Layer 1 Adaptation (functional module)
L1F	Layer 1 Forwarding (functional module)
L2A	Layer 2 Adaptation (functional module)
L2F	Layer 2 Forwarding (functional module)
L2SC	Layer 2 Session Control (functional module)
L3F	Layer 3 Forwarding (functional module)
L3RC	Layer 3 Resource Control (functional module)
L3SC	Layer 3 Session Control (functional module)
LER	Label Edge Router
LSR	Label Switching Router
MAT	MAC Address Translation
MEP	Maintenance Association End Point
MSBN	Multi-Service Broadband Network
NAT	Network Address Translation
NIC	Network Interface Card
NNI	Network to Network Interface
NTE	Network Termination Equipment
PC	Policy Control
POS	Packet Over SONET/SDH
PPP	Point-to-Point Protocol
QoS	Quality of Service
RG	Routing Gateway or Residential Gateway
SI-NNI	Service Interworking NNI
STF	Scheduling and Traffic Filtering (functional module)
TE	Traffic Engineering
TR	Technical Report
UNI	User to Network Interface
VMAC	Virtual MAC
WIP	Work in Progress
WT	Working Text

3. Technical Report Impact

3.1 Energy Efficiency

TR-145 is a converged framework to describe a convergence network that can support a multiplicity of residential and business services over a common infrastructure such that fewer network elements are needed; therefore energy consumption is expected to be lower than when deploying and operating multiple service-specific networks next to each other.

3.2 IPv6

TR-145 needs to support both IPv4 and IPv6 equally. Since TR-145 focuses on the data plane and some of the signaling (such as multicast), IPv6 needs to be supported in the appropriate functional Modules
The evolution to IPv6 can be step by step depending on the service providers' needs. TR-242 [16] describes the various IPv6 migration routes in detail.

As TR-145 also supports the service set from TR-101 [3] , all the IPv6 work that builds on TR-101 (TR-177 [8], TR-187 [9]), will be supported through TR-145 as well.

3.3 Security

Enhanced security is needed largely because of the network convergence, in particular the transport of business services and support for mobile backhaul, in addition to residential services. TR-145 security requirements are described in Section 4.6.2.

3.4 Privacy

Privacy involves the need to ensure that information to, from and between customers can only be accessed by those who have the right to do so. This of course primarily means limiting (information) access to the intended customers themselves, but there remains a general need for lawful intercept. While this document does not define any specific mechanisms to support such intercept, it ensures that it is not precluded.

There are 2 ways to ensure privacy, preventing data from being copied to a non-intended destination, and encrypting the data so that it cannot be understood even if it is intercepted. This document includes various mechanisms to prevent data from being diverted, e.g. not allowing hair-pinning between users at an access node, but content encryption is normally done end to end, and at the application layer and so beyond the scope of this document. There is a significant overlap between security and privacy.

4. Reference Network Architecture and Network Requirements

This section lists the technical challenges posed by the business requirements of TR-144 [4]

4.1 High Level Reference Model

This section describes a high level functional reference model of the Regional and Access networks. It consists of "Customer Location Function Set", an "Access Function Set", and "Aggregation Function Set", as shown in Figure 2.

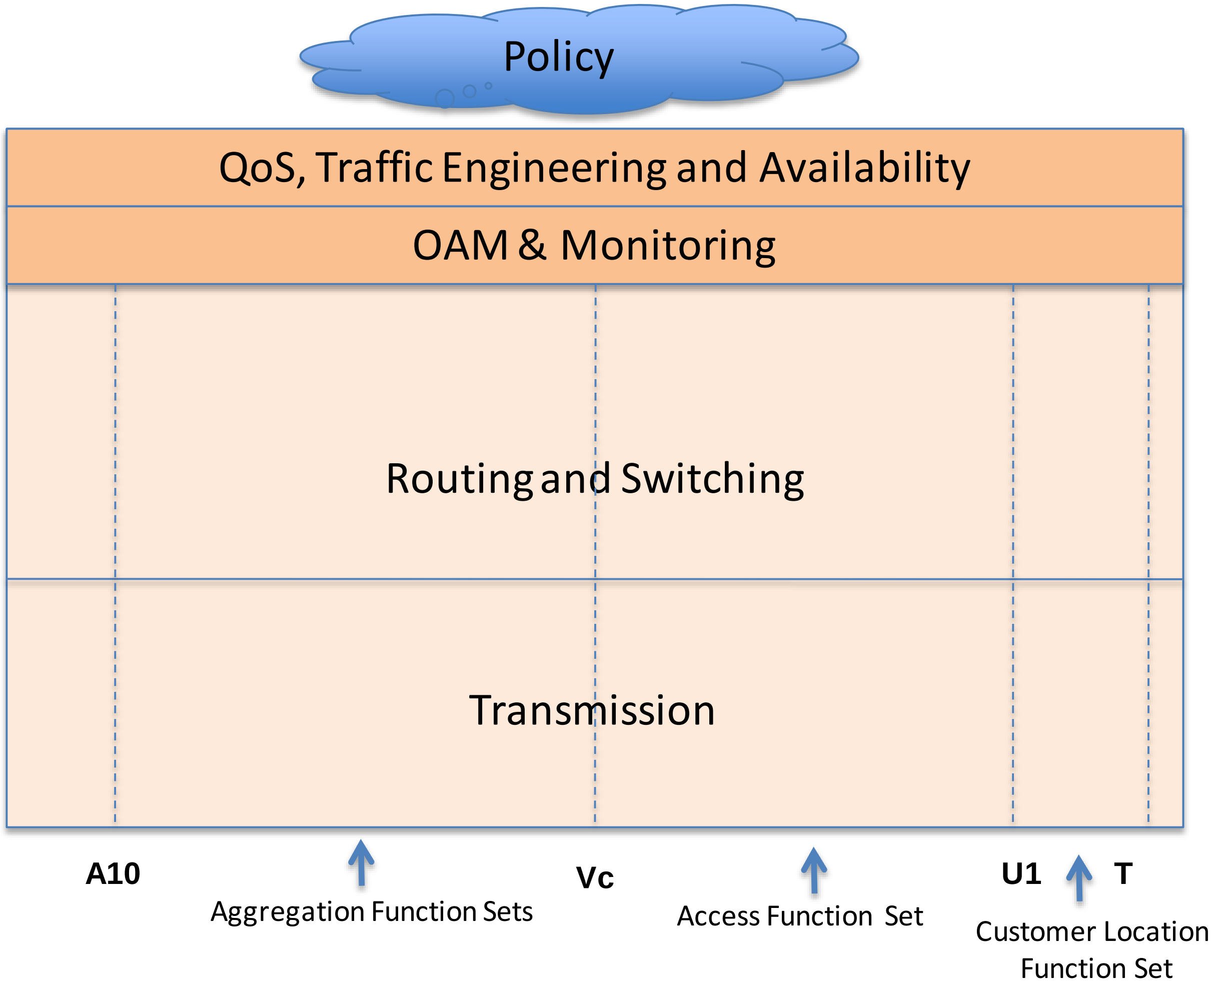
Policy
QoS, Traffic Engineering and Availability
OAM & Monitoring
Routing and Switching
Transmission
A10
Aggregation Function Sets
Vc
Access Function Set
U1
T
Customer Location Function Set

Figure 2 End to End Capabilities

4.2 Reference Points

Note: not all Reference Points in this section are represented in Figure 2

4.2.1 T

The T reference point is between the end user equipment and the network provider equipment.

For instance, for a carrier Ethernet service, it is between the end user and Ethernet service provider domains. For mobile backhaul, it is between the mobile carrier and the mobile backhaul service provider domain and for broadband access services, it can either be between the RG and other CPE in the Customer Location or between a B-NT and an RG.

4.2.2 Vc

Vc is the reference point at which the access functional set and the aggregation functional set(s) interconnect. This reference point is likely to be internal to the Access Node i.e. it is usually not instantiated as a physical interface.

4.2.3 Va

Va is the reference point at which a first level of Ethernet aggregation and the rest of the network interconnects. It may or may not be external to the Access Node. It can instantiate logical interfaces such as an I-NNI and/or can instantiate business interfaces such as an E-NNI-L2 (e.g. wholesale handoff in a distributed manner). For an explanation of logical interfaces such as an I-NNI, see Section 4.3.2. The Va reference point is shown in Figure 5 (see Section 4.4).
In WT-178, an Access Node with an internal Va reference point will use the V reference point for its uplinks.

4.2.4 U1

The U1 is the reference Point at which the Access Function Set interconnects with the Customer Location Function set, (e.g. Between xDSL modem (xTU-R) and RG function for consumer services). This Technical Report separates U1 interfaces into one of two categories: packet based interfaces, such as IP/Ethernet, and transport based interfaces, such as DS1.

Note that TR-145 uses the U**1** reference point, which is different from the U used previously in TR-101. U1 and U are two distinct reference points. The TR-101 U reference point instantiates the access technology, e.g. xDSL.

4.2.5 A10

A10 is the reference point at which the Multi-Service Broadband Network and ASPs/NSPs' POPs interconnect and is typically the reference point that handles more aggregated business handoffs through the E-NNI interface

4.2.6 U

In the TR-145 architecture the U is within the access function set (that has been defined by the ITU), and is therefore usually invisible in this document. U is only addressed in the access link technology specific section (Section I.1). U will be an exposed interface in the nodal specifications defined in WT-178. Therefore every interface instantiated at U1 in this document can be instantiated at the U reference point in WT-178.

4.3 Interfaces Definitions

4.3.1 Business Interfaces

This section defines enhancements and capabilities of the following business interfaces to support the multi-service applications.

- E-NNI for application/service provider handoff
- UNIs for customer hand-off

Figure 3 depicts the business interfaces and their relationships to the reference points.

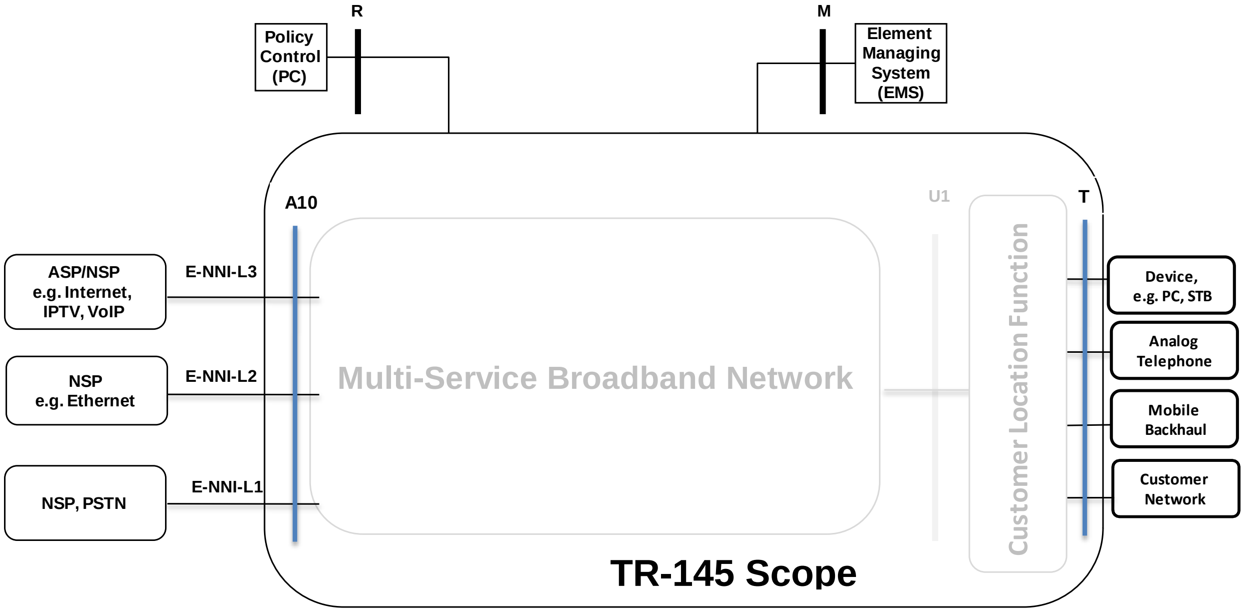

Figure 3 External Business Interfaces and Reference Points

4.3.1.1 The E-NNI Business Interfaces

The E-NNI serves as the business interface between the regional/access network and ASP and/or NSP as specified in TR-101 and TR-144 [4]. There are several classes of E-NNI interfaces, such as E-NNI-L1, E-NNI-L2, and E-NNI-L3.

E-NNI interface can be Layer 2 or 3 and supports the following types of interface:

- Ethernet (layer 2)
- IPoE/PPPoE (layer 3)

E-NNI-L1 interface is at the transport layer supporting:

- Circuit

- Wavelength

The E-NNI business interface can be instantiated at various reference points in the architecture, depending on business requirements e.g. handoffs at Va and/or A10, The UNI Interfaces

The UNI can be concatenated and instantiated at the U1 (named UNI_2) or at T (named UNI@U1) reference points depending on the service, e.g.:

- For Residential Triple-Play, the UNI can be at U1 or at T (depending on whether RG is provided or not)
- For L2VPN Services, the UNI can be at U1 (handoff between Ethernet Service Provider and Regional Broadband Provider) or at T (hand-off between end customer and Ethernet Service Provider). So depending on who is providing the service to whom, there can be several UNIs.
- For a ‘transport-only’ service , the UNI is at U1

The UNI interface at the T reference point includes:

- Ethernet over any physical home network technologies (UNI-L2)
- POTS/ISDN (UNI-L1)
- TDM T1/E1 (UNI-L1)

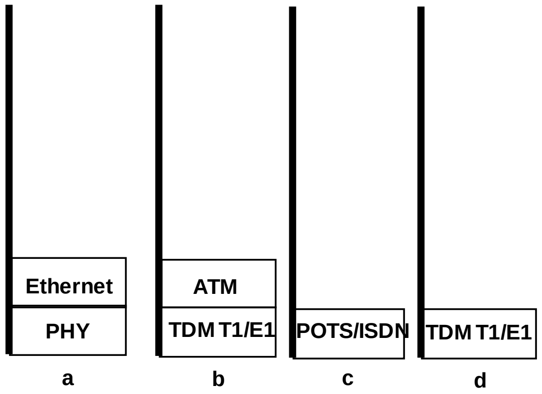

Figure 4 Protocol Stacks at UNI interface

4.3.2 Logical interfaces definitions

A logical interface in this document is defined as a boundary between two functional modules. While some logical interfaces are also physical ones, not all logical interfaces need to be instantiated as a physical interface. Several functional modules may be grouped into a single physical node. In such a case, the logical interfaces internal to the node will not be exposed
The business interfaces, such as E-NNIs and UNIs defined in Section 4.3 are logical as well as physical interfaces.

This section defines additional logical interfaces:

- ◊ **Internal Network to Network Interface (I-NNI)** between two functional modules. The I-NNI will implement the same protocols and encapsulation technology on either side of the interface. Note that internal here refers to 'internal' to a service provider network rather than internal to a node.
- ◊ **An Adaptation NNI (A-NNI)** is between an adaptation module of a given technology type with a functional module of a different technology type. For example MPLS encapsulation/tunneling of Ethernet frames.
- ◊ **A Service Interworking NNI (SI-NNI)** that interconnects two different functional modules which are not at the same level of the OSI stack e.g. interconnecting an Ethernet functional module with an IP aggregation functional module.

4.3.3 Applicability of interfaces to reference points

The following table shows how the various interface types (logical and/or business interfaces) map to the various Reference Points defined.

Reference Point Interface Type	**T**	**U/U1**	**Va**	**A10**
UNI	Yes (e.g. customer equipment to NID	Yes (e.g. RG to Access node)	N/A	N/A
I-NNI	N/A	Yes (e.g NID to Access node	Yes	N/A
E-NNI	N/A	N/A	Yes	Yes

Table 1 Applicability of interfaces to reference points

The A-NNI and SI-NNI interfaces are not included in the table as they will never be instantiated at a known reference point and are always internal to a node. The Vc reference point is also not included in the table as no interfaces are instantiated on it.

4.4 Detailed Reference Models

Figure 5 shows how aggregation functionality is divided into three distinct functional sets, namely:

- IP Service Layer: these are the services seen by subscribers at the IP layer e.g. L3-VPNs or Residential Internet Access.
- The Ethernet Service Layer: these are the services seen by subscribers as Ethernet Services (MEF Services) as well as the (emulated) Ethernet constructs to connect IP RGs or CPE to their IP Service Edges. The Ethernet Service Layer will perform 802.1ad Ethernet Aggregation, very much like the model deployed in TR-101.
- The Supporting Aggregation Layer: this is the layer that emulates the Ethernet Service layer on top of a different technology (e.g. MPLS).

Figure 5 also shows the business interfaces UNI, which can be either @U1 or @T (hence the Ethernet Service Layer terminates at U1 or T respectively), as well as the E-NNI-L2 and E-NNI-L3 business interfaces @A10. The UNI can have L2 and L1 instantiations. The L2 variant is called UNI or UNI-L2, while the L1 variant is denoted as UNI-L1.

The logical connectivity between the IP Service Layer and the Ethernet Service Layer is across the SI-NNI logical interface, and the adaptation of Ethernet onto the Supporting Aggregation Layer technology is via the A-NNI.

The Underlying L1 Transmission Infrastructure is connecting to the Ethernet and Supporting Aggregation layers through A-NNIs and has business interfaces (E-NNI-L1 and UNI-L1) as well.

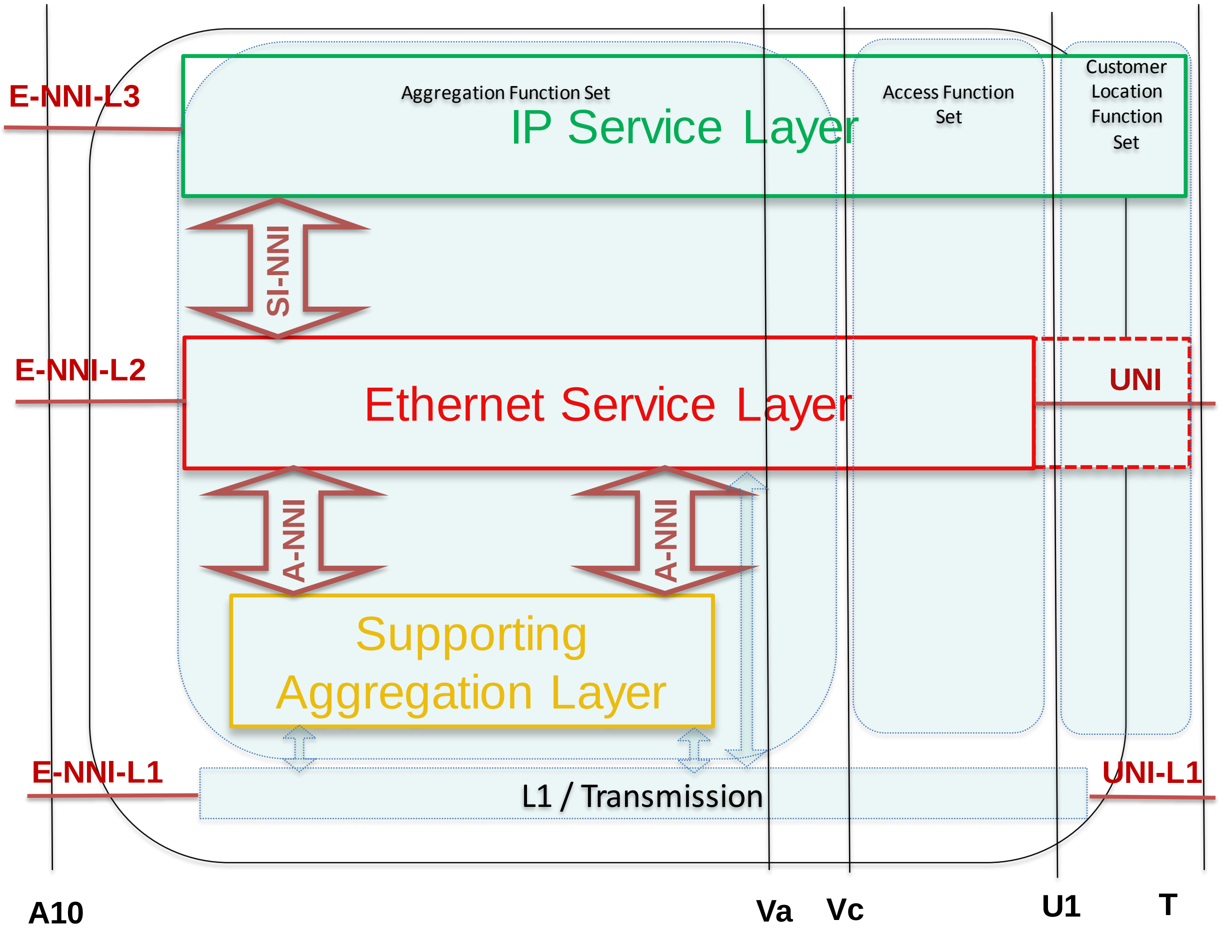

Figure 5 Detailed Reference Model of Multi-Service Broadband Network (MSBN)

The Access Function Set refers to the set of packet based access technologies (e.g. various DSL technologies, PON, Ethernet) as well as transport specific access (e.g. SONET, WDM, etc.) that can be used to support the various services specified in TR-144 [4]. Since these access technologies have been already been defined by ITU and other SDOs, we simply recognize various access technologies and abstract them as an "Access Function Set".

The Customer Location Function Set provides functions, such as bridging, routing, and subscriber session control for consumer, service handoff to mobile carriers, as well as the Ethernet service termination function.

4.5 Infrastructure Virtual Connections

An Infrastructure Virtual Circuit (IVC) is a building block for constructing end to end L2 service connectivity. An IVC can be from either user to user or from user to service edge. An IVC is the association between logical or business interfaces e.g the association between the UNI and the SI-NNI. Another example is the association between the ENNI-L2 and the UNI. Both examples are depicted in Figure 6 as IVC1 and IVC2 respectively.

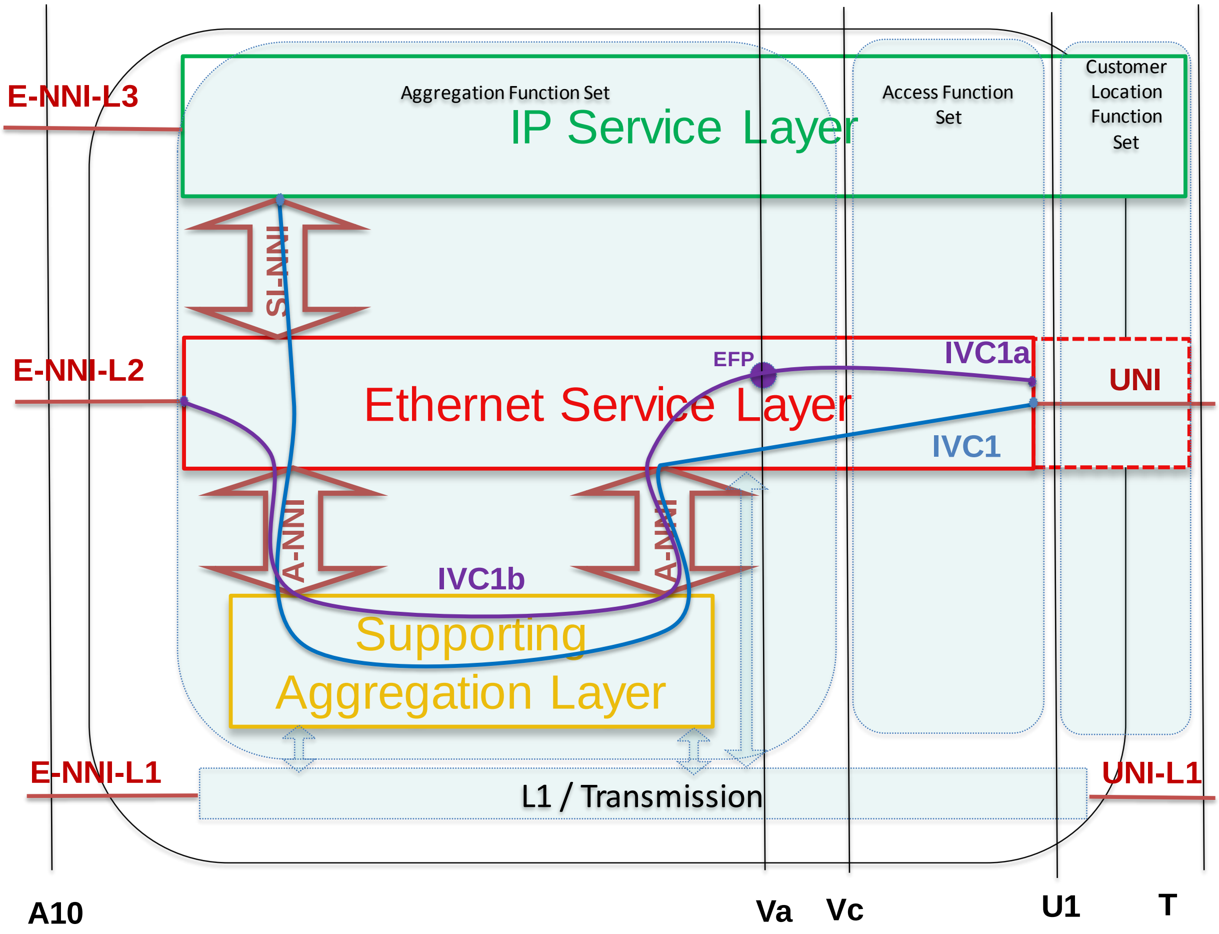

Figure 6 IVC and IVC segment Example

Note the IVC terminology here uses a similar philosophy to the one used in MEF, but there is no direct mapping between a MEF OVC and a BBF IVC. The MEF OVC is a construct to create end to end Ethernet Virtual Connections between MEF UNIs across different Service Providers. It is therefore an association between a UNI and a MEF E-NNI. The BBF IVC meanwhile is a construct that can have associations between a greater set of Interface types such as UNI to UNI, UNI to SI-NNI, UNI to ENNI-L2. The BBF IVC can be seen as a superset of the MEF OVC.

An end to end IVC consists of one or more IVC segments. For instance, an end to end multipoint IVC can be built from the combination of a simple pt-pt (E-LINE) IVC segment in the access network (IVC1a in Figure 6) and a multipoint IVC (IVC1b in Figure 6) in the Regional network that supports E-LINE, E-LAN, and E-TREE services. This example IVC extends the E-LINE, E-LAN, and E-TREE services all the way to the end user.

A pair of IVCs can be used to provide redundancy, which is especially important for the RNC/GW in mobile network and core network for mobile backhaul. This can be achieved with or without geographical diversity in the network

A particular logical interface could be shared by multiple virtual connections; therefore there is a need distinguish a single virtual connection, such as an IVC on a logical interface. The IVC is an end to end construct. The instantiation of an IVC at a given logical interface is called the **Ethernet Flow Point (or EFP**, see Figure 6 for an example**)**. By defining the functions and attributes at the EFP level, one can model how packets traverse a certain function and get encapsulated/conditioned/ filtered/etc.

4.6 Network requirements derived from service definitions

The abstract reference architecture includes end to end capabilities, like policy, QoS, and OAM, in the broadband network as shown in Figure 2 .

The rest of this sub-section details these end to end requirements.

Note: the term ‘aggregation node’ refers to a node that is purely performing aggregation functionality for a given network aspect. Aggregation can be performed at the IP layer, or at Layer 2 (Ethernet/MPLS). The term ‘access node’ refers to a node that is performing access functionality but also provides the first or only aggregation function.

4.6.1 Resilience

The requirement for a multi-service access and aggregation network to provide mobile backhaul as well as to provide critical business and residential services to fixed-line customers drives the requirement for high availability of network connectivity.

Various levels of resiliency can be achieved by choosing an appropriate topology, as described in Appendix II.9, as well as by providing intra-node functionality.

Dual homing an access node to a single aggregation node provides service resilience in the case of link failure or a network card failure. Moreover, dual homing of an access node to two different aggregation nodes or IP aggregation nodes provides service resilience in the event of an aggregation node outage (due to failure or routine maintenance). This is referred to as a dual-homed device (DHD). Network based resiliency mechanisms such as dual homing should not change the service offering (apart from increasing the availability). For example, if the end to end service is an E-Line, dual homing the access could result in E-LANs or E-TREEs being produced.

A pair of aggregation nodes might be needed to dual-home a ring of access nodes (Dual Homed Network or DHN). Measures need to be taken in order to avoid blackholing traffic as a result of a break within the ring.

The availability of some real-time services offered to customers is expected to be 99.999%, which is equivalent to an unavailability of only 5 minutes per year or less than 1 second per day. Further, if it is desired to make network events truly imperceptible from the customer's point of view, depending on the application this could require each interruption to last no more than 50 ms.

Hardware redundancy, software redundancy, link, node and path (physical or virtual) redundancy, as well as state mirroring, are methods that can be used to help meet the resiliency goals. At the link level there is a need to have a tight coupling between the packet layer and the underlying layer(s), including leveraging OAM events. Additionally the application can add more information to the data streams to cater for momentary loss, an example being Forward Error Correction for IPTV traffic. Retransmission at the application layer or physical layer is other examples. In the extreme case the application can duplicate streams across different network paths such that no loss is perceived by the receiving terminals. The network elements are then required to select the preferred stream on the fly.

Network convergence times are driven by service requirements. The network architecture must be able to use various resilience techniques such as the ones above, on a per service basis, where possible.

As a foundation, link and node resiliency methods should be built into the architecture and the control plane, and link and node failures should result in a maximum loss of connectivity of a 1 second duration. On top of that, (virtual) path redundancy mechanisms can be deployed to ensure sub-50 ms convergence for certain services.

The broadband network is also expected to ensure continuous availability between the RGs of broadband customers and the IP network of the Broadband Service Providers, even when users are not sending any traffic over their broadband connection. This continuous connectivity is to support for example supervision, firmware upgrades, incoming calls etc.

4.6.2 Security

The broadband network is expected to provide sufficient subscriber isolation and service integrity to prevent any individual subscriber initiating a denial of service attack, or theft of service from the network or another customer. The broadband network must also prevent inappropriate leaking of information or content between subscribers.

Various mechanisms are defined to protect the network against distributed denial of service attacks (DDOS = a very large number of customers are subverted and synchronized to launch a given attack simultaneously), e.g. by rate limiting various types of control packets, but it is extremely difficult to fully protect a network against the wide range of such possible attacks. Another is the support of Access Control (e.g. by specify lists or filters).

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.

For a business service, then network availability is another key aspect of security. Support is provided for various topologies that can provide duplicate or back-up paths, and therefore improve availability. Quality of Service mechanisms can also act as a security by allowing more important (e.g. business or control) traffic to get through when a network is congested.

End-user security is also important and has a strong overlap with privacy. Again L2 separation of customer traffic and preventing hairpinning are among the main techniques covered in this document. There is also support for tunneling. Encryption remains 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 and so beyond the scope of this document, but care has been taken to ensure that these techniques can be used transparently across the networks described.

4.6.3 Quality of Service and Resource Control

Since an architecture that meets the requirements of TR-144 [4] needs to support different services concurrently, QoS features therefore need to be supported as part of the functionality. However, the functionality will be different, both in terms of features, as well as scale, depending on its location in the network.

The architecture needs to support multi-edge deployments. This presents new challenges with regard to QoS and resource control.

The architecture needs to support at least the following QoS functions regardless of the place in the network:

- DiffServ queuing per service class inside aggregation (L2 and/or L3) nodes in order to support multi-edge (this was already a TR-101 requirement)
- Aggregation functionality (L2 and/or L3) needs to be able to 'divide' bandwidth into chunks (egress function) with Diffserv per service/class queuing within those chunks. This is often called 'bandwidth partitioning'
- Aggregation functionality (L2 and/or L3) needs to be able to protect the aggregation and core network (ingress function)
- The aggregation functionality (L2 and/or L3) controls how much bandwidth leaves and enters the network, how it is treated, and how classes can use other classes unused bandwidth

Some network services use shaping and policing to make sure that the network administrators can control how much traffic leaves or enters the network. Within that shaped/policed bandwidth, per subscriber queuing could be leveraged to offer fairness. Note however that the impact of shaping and policing will be service dependent.

The Multi-Service Network Architecture 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 and admission control may be needed). For example walled garden type applications like Voice and Video could bypass per subscriber scheduling in a multi-edge environment. As these are delay sensitive (voice) and/or loss sensitive (video VoD), the network administrator has to estimate what bandwidth is needed on every link to ensure these applications work with minimal delay and loss. As we cannot use loss due to congestion as a way to slow down voice and video frames, we have to make sure that all voice and video flows are controlled and accounted for in the architecture. A video or voice flow that would exceed the available capacity should not be allowed to start. In order for this to happen, some form of resource control is needed.

Resource control can be Off-Path and On-Path. Off-Path resource control involves a centralized system that has awareness via various out-of-band means to the current state of the network. There may be an interaction between the Application Layer and this system to determine if the application is allowed the required network access.

On-Path resource control refers to the application layer requesting network resources by in-band signaling. This request will either be accepted or denied.

Resource control is not always needed, e.g. when the network has appropriate capacity dimensioning, so that all managed voice and video calls can always be carried.

One particular problem is ‘unified multicast and unicast resource control’. Not all Multicast replication points are resource control capable or indeed need to be. A resource control system may need to be able to take into account both unicast streams and multicast streams on a per subscriber basis, in order to make correct resource control decisions.

Service Edges can still implement per customer QoS in addition to the above.

There is also a need to adapt the aggregate speed of Ethernet interfaces at the A10, U1 and Va reference points to rates lower than the line-speed, e.g in order to interwork with underlying infrastructures), or to support hand-off at arbitrarily defined rates.).

4.6.4 Policy control, and AAA

TR-145 focuses on the data plane in the access and aggregation network. Functional modules performing layer 2 and layer 3 session control defined in later Sections provide the interaction with the Policy Control System (PC) via the R reference pointto perform the following functions:

- Session Authentication, Admission Control, and establishment/termination
- configuring QoS scheduling and conditional access control attributes

In TR-134, a functional architecture is defined for Policy control and AAA. TR-134 distributes the above functionality between the R reference point for policy control, and the B reference point for AAA.

However, in this document, policy control and AAA are both performed via the R reference point.

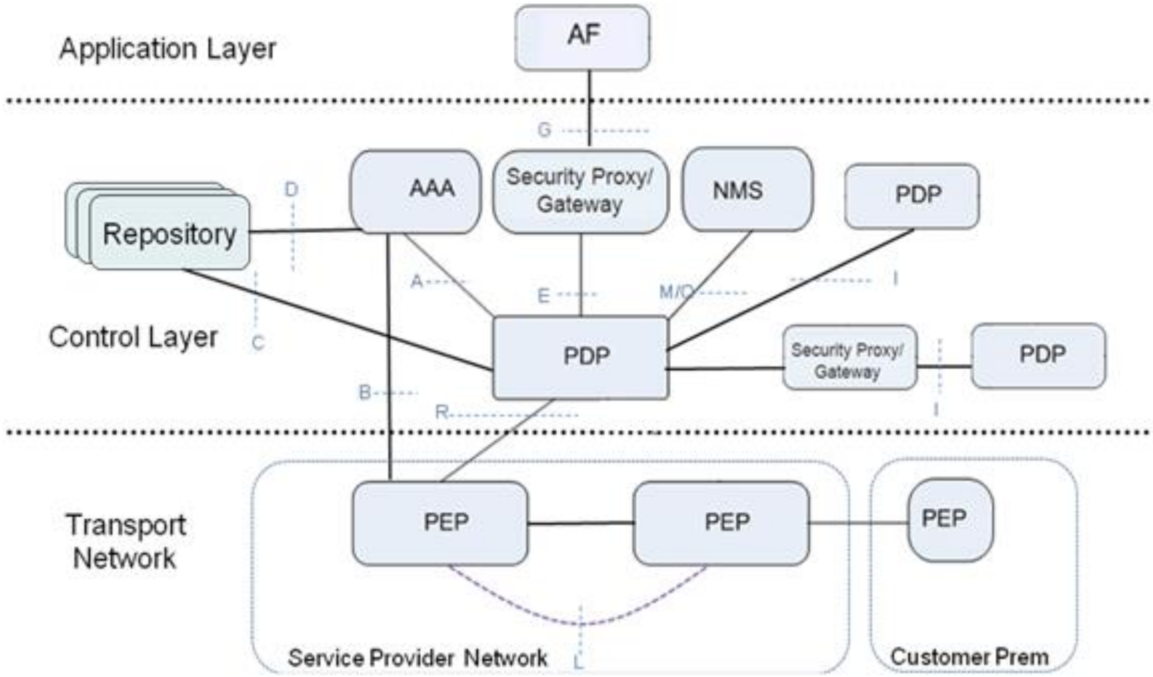

Figure 7 BPC Framework Interface Architecture

Please see TR-134 for further information.

4.6.5 IP Multicast

Multicast can improve network efficiency by sending the same IP traffic to a group of receivers in a single transmission across a dynamically signaled multicast forwarding tree that is set up as a result of end-user signaling. TR-101 already described 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 architecture needs to support and extend the TR-101 functionality to support multicast delivery across IP-VPNs for business customers and to support multicast wholesale services at L2 and L3 utilizing network virtualization. As in TR-101, the architecture needs to support both Any Source Multicast (ASM) and Source Specific Multicast (SSM). Lastly, the architecture needs to support both IPv4 and IPv6 multicast.

For the multicast services that need resource control, multicast resource control (i.e. controlling the amount of multicast traffic that can be replicated across a specific interface) has to be supported.

4.6.6 Subscriber Management

The Subscriber management function encompasses the following sub-functions:

- Session identification and access control – (Sections 5.1.2.2 and 5.1.2.3)
- QoS scheduling – that includes fine grain QoS scheduling on L2 or L3, policing and shaping (Section 5.1.3)
- policy control functions - (Sections 5.1.2.2 and 5.1.2.3)

4.6.7 Network management

The modules defined within this document may be interfaced and interact with a Network management system, but this is not defined within this document.

4.6.8 OAM and Interface Management

The architecture needs to support the ability for each provider to 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. If different levels of virtualizations are used (e.g. VPLS over MPLS pseudowires over Ethernet links over WDM), every layer needs to have its own set of OAM functionality. Alarms per layer must be generated. Consequent actions on upper layers may be enabled.

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 hand-off points (UNI, E-NNI), there must 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).

4.7 Topologies

Star, ring, and subtended network topologies, as well as any combinations thereof as described in Annex B/TR-144 [4], are intended to be supported by TR-145. This section gives more details on the topology in access and aggregation network based on generic network deployment trends, redundancy needs, and the logical connectivity perspective.

Access network topologies:

In order to meet the demand of increasing bandwidth ANs are being located closer to the end-user. Therefore the following interconnect topologies need to be supported, possibly in combination:

- Point to multi-point e.g. PON-fed ANs,
- daisy chained ANs
- Hub and Spoke.
- Ring Based

ANs themselves are becoming more **distributed** i.e. they may have subtended components such as remote shelves, or sealed modules.

Driven by **redundancy** considerations, the access network will need to support:

- Rings of access nodes (to improve overall availability with minimal extra trunking)
- Meshing (Dual homing to the same or a different AN – the latter is primarily to support business customers who need full redundancy)

From a **logical connectivity** perspective, the access network needs to support:

- Point to Point: e.g. E-Line
- (Rooted) Point to Multipoint e.g. E-Tree
- Multipoint to Multipoint e.g. E-LAN
- Multi-layer topologies (L1, L2, L3), where L1 is emulated over a packet L2/L3 topology.

Aggregation network topologies:

All the physical topologies described in the access network section also apply to the aggregation network.

In addition logical multi-rooted E-TREE topologies need to be supported e.g. for multiple service edges. Note that the logical topologies may be layer dependent.

Driven by **redundancy** considerations, the aggregation network may also need to support:

- Dual homing of Access Nodes. This requires additional control and signaling for load sharing and/or switch-over, and this applies to both dual-homed device and dual-homed network configurations.
- Redundant Service Edges including scenarios where they are not directly connected to the Access Node. This will require control/signaling for switchover and fast re-route.

From a **logical connectivity** perspective, the aggregation network needs to support the same characteristics as listed in the access network section but also the appropriate mechanisms to achieve connectivity to redundant Service Edges.

5. Functional modules specifications

This section provides a description of the network functional modules, technology independent functional reference model, use cases, and the requirements organized per functional module. This section does go into technology to the degree necessary, but it remains agnostic as to which node needs to implement which technology.

5.1 Technology Independent Functional modules

5.1.1 Layering Functional Module

This section adopts a simplified layering definition that separates the associated functions for a given layer into two classes of function: Adaptation/termination and forwarding.

Adaptation and termination encompass all the functions associated with the mapping of client signals onto a given layer and the associated layer end-point processing functions such as encapsulation/decapsulation and OAM flow termination.

The **Forwarding** function includes the relaying of layer information, and any intermediate point OAM processing.

The layer naming convention acknowledges the OSI convention of physical, link and network layers: 1, 2 and 3 respectively, but the use cases will acknowledge that recursion within and across the layers has become commonplace. For example, an implementation of TR-101 (which uses 802.1ad constructs) onto 802.1ah onto MPLS-VPLS would see several hierarchical Adaptation steps within what is commonly considered to be a layer 2 technology.

5.1.1.1 Layer 1 Adaptation (L1A)

The function that adapts client signals onto the physical layer technology and provides the associated termination function at the physical layer. This may also involve mapping of higher layer information into the layer 1 information structure for transmission (e.g. G.709 Digital Wrapper adapting packet onto wavelength services).

5.1.1.2 Layer 1 Forwarding (L1F)

The function covers relaying/switching function for layer 1 technologies such as SONET/SDH. The L1F function is commonly termed as a digital cross connect or "time" switch. However L1F functionality will not be discussed in any detail in this document as its behavior has no impact on the transported packet layers.

5.1.1.3 Layer 2 Adaptation (L2A)

Is the function that adapts and maps client signals onto layer 2 including the resolution of the forwarding information to reach the peer L2A entities, encapsulation, and the associated termination on layer 2. Examples of resolving the address of the peer L2A entity include IP to MAC resolution via ARP and C-MAC->B-MAC lookup per IEEE802.1Q-2011. The Adaptation could be done based on layer 2 or higher layer information. This can recurse within the layer (e.g. stacking of an S-tag on a C-tagged Ethernet frame, VLAN mapping into a Pseudowire). The use of multiple tags (stacking) should be supported.

5.1.1.4 Layer 2 Forwarding (L2F)

A frame switching function for the given layer 2 technology that must be capable of selectively switching or multiplexing to offer p2p, p2mp or mp2mp connectivity. L2F includes the Multicast Replication Function (MRF) that allows replicating the packets of a multicast group to dedicated leaves based on layer 2 (MRF-L2).

5.1.1.5 Layer 3 Forwarding (L3F)

An IP packet forwarding function with longest match lookup. It also encompasses the full set of intermediate node functions for the layer (e.g. for IPv4: ICMP etc.). L3F includes the Multicast Replication Function (MRF) that allows replicating multicast packets of a multicast group to dedicated leaves based on layer 3 (MRF-L3).

5.1.1.6 Legacy Interworking Function (LAF)

A function that performs adaptation from legacy protocol/interfaces to a packet interface.
TR-144 [4] defined legacy services including POTS, TDM (T1/E1 TDM) and ATM over TDM.

5.1.2 Control Functional Modules

5.1.2.1 Session Control Overview

Session control correlates different traffic streams into a session (which can be L3, L2 IVC, or PPPoE coming from the same UNI interface) based upon gleaning of associated control plane and/or data plane traffic within those traffic streams. It may support authentication, admission control, session termination and other functions. There can be an interface to the Policy Control (PC) function in order to build the mapping table and configuring QoS scheduling and conditional access control attributes.

5.1.2.2 Layer 2 Session Control function (L2SC):

This is the function that provides the control of L2 session establishment and termination. A L2 session can correlate different layer 2 traffic flows based upon gleaning of associated control plane and/or data plane traffic (e.g. VLAN id's) within those traffic flows. It may support authentication, L2 label assignment (e.g. VLAN id), IVC stitching, Configuring QoS scheduling, Admission Control, provisioning of L2 OAM, and may or may not involve communication with an external AAA and/or Policy Control System.   The L2SC can be centralized or distributed but must be able to interact with L2A and L2F functions.

5.1.2.3 Layer 3 Session control function (L3SC):

The control of L3 Sessions is dynamic and based upon the gleaning of associated control plane and/or data plane traffic, the different aspects of L3SC are defined by the following general areas and the associated requirements.

- The L3SC Function is presented with data packets by the L3F Function, and then applies the necessary context to compose the data streams into L3 Sessions.
- The L3SC Function instantiates the output from data plane policies in the STF Function and manages these for the duration of the session.
- The L3SC Function interfaces with the Policy Control System, which may be used to manage attributes of L3SC. Other methods of managing the L3SC include normal network management such as CLI/SNMP/NMS/OSS.

5.1.2.4 Layer 3 Resource Control (L3RC) Function

In a multiservice network, the variability of the traffic over time, and the possibility of loss due to congestion can be quite high, especially as the take up rate of unicast video or videoconferencing is hard to predict. One method to solve this is by doing resource reservation or resource availability checking, i.e. the applications interrogate the network infrastructure to see if there is sufficient available capacity before starting to stream the video traffic. Multicast based video services can also suffer from similar problems. For the same reason outlined above, loss due to congestion should be avoided within the network infrastructure for these service types. When traffic is multicast, a small period of congestion can lead multiple users to have a badly impacted quality of experience. Any multicast replication point therefore should be capable of keeping track of the amount of already replicated multicast traffic on any given interface, and should prevent the admission of additional multicast groups that could cause congestion within the video class.

The L3RC functional module handles resource reservation for both unicast and multicast traffic. The L3RC functionality must be tightly coupled with the L3SC and STF functionality, as the L3RC functionality depends on state within the L3SC and STF functionality.

5.1.3 Scheduling and Traffic Filtering function (STF):

This function has two elements:

- Scheduling, which can be simple or hierarchical, is applied at layer 2 or 3. This may involve queuing, shaping, and policing. This provides the proper treatment to traffic depending on its QoS classification, both on ingress and egress. The traffic treatment applies to data plane, control plane and the management plane traffic. Classification functionality is regarded as being part of the SF in this document.

- Traffic Filtering, which includes conditional access control (e.g. ACL filtering). For multicast, TF performs Multicast Group Control Function (MGCF) that enables the joining and leaving of multicast groups. This can be configured locally through black/white list entries or controlled remotely based on AAA/policy interaction.

5.1.4 Synchronization function (Synch):

The function that extracts and can relay high quality frequency and/or phase and/or time synchronization, to allow a network to support applications requiring such synchronization. These applications include mobile backhaul, circuit emulation, IP delay monitoring, alarm/ / performance time-stamping. The relevant functionality is defined in G.810 [[26] and G.8260 [25].

5.2 Multicast Architecture

Various flavors of Multicast implementations related to layer2 and/or layer 3 technologies exist. The choice of L2 versus L3 techniques for multicast is often governed by the network location of the replication points, as well as the service type they aim to support. L3 multicast forwarding techniques can be leveraged to optimize forwarding across arbitrary topologies whereas L2 multicast forwarding can be used where the topology is simpler.

Multicast can be decomposed into three functions - multicast replication, replication control and multicast group control. These sub-functions are nested in the more generic functions modules, such as, L2F, L3F, L2A, and STF as described in the above sections.

The control relationship among the multicast sub-functions above is shown in Figure 8. The Multicast Group Control Function can be controlled by an external network element (e.g. TR-147 sec 6.4, AAA, EMS, etc.), or can be statically configured. The Multicast Group Control Function controls the Multicast Replication Control function, which controls the Multicast Replication function.

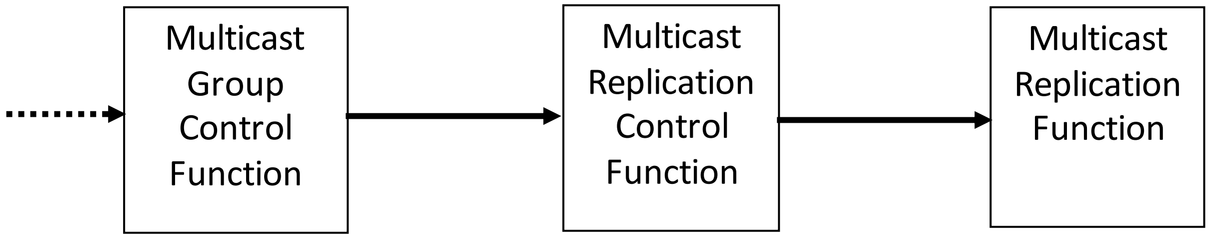
Multicast Group Control Function
Multicast Replication Control Function
Multicast Replication Function

Figure 8 Control relationships between multicast sub-functions

The Multicast Replication function should be positioned at all multiplex points in the network (starting as close as possible to the customer, but taking the deployment economics into account). This network design principle is independent of layering. The Multicast Replication Control Function is closely tied to the Multicast Replication Function (e.g. IGMP snooping). Multicast stream availability needs to be able to be filtered on a per subscriber basis (white lists/black lists) by the MGCF.

TR-145 supports multicast for both IPv4 and IPv6. Multicast Listener Discovery (MLD) is the IPv6 equivalent of IGMP in IPv4. The basic functionality of intercepting MLD packets, and building membership lists and multicast router lists is the same as for IGMP. It has two versions: MLDv1 is similar to IGMPv2; MLDv2 is similar to IGMPv3. MLDv1 is defined in RFC 2710 and MLDv2 is defined in RFC 3810. Both MLD versions support the Any-Source Multicast (ASM) model. MLDv2 supports the Source-Specific Multicast (SSM) model, whereas MLDv1 supports the SSM model only, with the help of SSM Mapping. MLDv2 uses ICMPv6 message types instead of IGMP message types.

5.2.1 L2 based Multicast Architecture

Multicast is implemented at various points in the service provider's network with specific functions depending on the technologies and network design. The design goal for L2 Ethernet based network is to support multicast optimization by controlling the flooding of Ethernet multicast frames using IGMP or MLD snooping [RFC4541], such that packets are replicated only on those ports (physical and logical) that have specifically requested a multicast group.

As stated in Section 4.6.5, the TR-145 architecture will continue to support TR-101 multicast. In addition it needs to be able to offer an extra level of virtualization in order to offer multicast for wholesale and business services as well. The virtualization can be implemented using L2 VLAN based architecture to offer multicast for various market segments, such as consumer, wholesale, and business. Specifically, one or more N:1 VLANs can be used to carry IGMP or MLD messages and multicast traffic for IPTV services. Another set of N:1 VLANs can be used for wholesale or business multicast services.

It is realized that for either wholesale or business multicast applications, it is possible that the customers are sparsely located especially within an access network, such that there might be fewer advantages of multicast replication within an access node. Therefore, in addition to the N:1 VLAN model, the combination of 1:1 VLAN in the access node and N:1 VLAN in the aggregation network, where the IGMP snooping and filtering function will be pulled back to, can work for some wholesale and business multicast. The architecture needs to support both these VLAN models to carry multicast data and signaling.

In L2 multicast, the data plane uses multicast MAC address filters which link L2 multicast groups to egress ports on bridging devices. In order to automate the setup of the filters, a bridge forwarding engine will redirect IGMP or MLD packets to MRCF-L2, the controlling function. Based on the requested IP multicast group, the bridge will set up a L2 multicast filter entry that allows or prevents packets to flow to the port on which it received the IGMP or MLD report.

Specifically, in the technology independent model, L2F will provide the MRF-L2 function, while L2A will provide MRCF-L2 IGMP or MLD snooping; and STF provides MGCF conditional access control, (e.g. ACL),.

The architecture must support multicast replication control on a per VLAN basis.

The detailed requirements of multicast can be found in Section 7.

Note, the mapping of ASM IP-Multicast addresses to a L2 address as defined in RFC 1112 may not be unique. 32 IP-Multicast group addresses can map to a single Ethernet multicast MAC. The architecture should provide equivalent multicast functionality for both IPv4 and IPv6. MLD multicast in a L2 multicast architecture needs to support both native Source-Specific Multicast (SSM), as well as the model based on source-unspecific MLD joins by using SSM Mapping in the network. MLD snooping [RFC4541] is used to intercept MLD packets at L2 device and set up a Layer 2 multicast forwarding table. This allows multicast capable L2 devices to forward multicast data to the ports that requested a multicast group.

5.2.1.1 L3 based Multicast Architecture

The architecture should provide equivalent multicast functionality for both IPv4 and IPv6. The architecture needs to support both Any Source Multicast (ASM) and Source Specific Multicast (SSM). Either can be implemented in access, regional or core networks and the choice is dependent on scalability and service provider's network architecture.

Using IGMPv3 or MLDv2 with L3 forwarding supports ASM and SSM:

1. ASM forwarding is based on the IP-Multicast address only which means that traffic from all senders in a group use the same multicast group address as in the packet's destination address, and these multicast packets are forwarded to all receivers.
2. Source Specific Multicast also uses the IP-Source Address as an input to the forwarding decision.

The Multicast Replication Control Function L3 (MRCF-L3) that controls joining / leaving multicast groups based on either IGMPv2 or IGMPv3 or MLD with Proxy Routing and/or multicast routing protocols in the IP network, is a sub-function of the L3F function described above. When IGMPv3 or MLDv2 is used, the MRCF-L3 will also control the replication of the multicast based on the source address of the multicast group according to source specific and any source features.

5.3 Technology independent functional reference model

5.3.1 Diagrammatic Convention

5.3.1.1 General:

The diagrammatic notation incorporated in this document represents technology independent modules as boxes, connected by lines representing interfaces. There is no convention as to where a line connects to a box and there is no implied hierarchy.

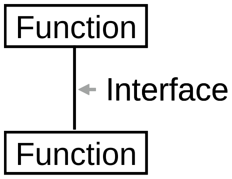

5.3.1.2 Nested functions:

Where functional decomposition is required, such functions are illustrated as being contained within the higher order function.

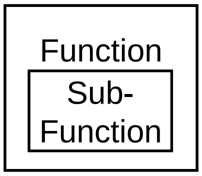

5.3.1.3 Interfaces and Reference points

Reference points are identified by a solid line. Interfaces are called out by a dashed line.

5.3.1.4 Layering specific Notes

Condensed notation
It is frequently not necessary to draw out all components of all layers, a shorthand would be to represent the set of lower layer interfaces and Adaptation functions as a dotted line interface.

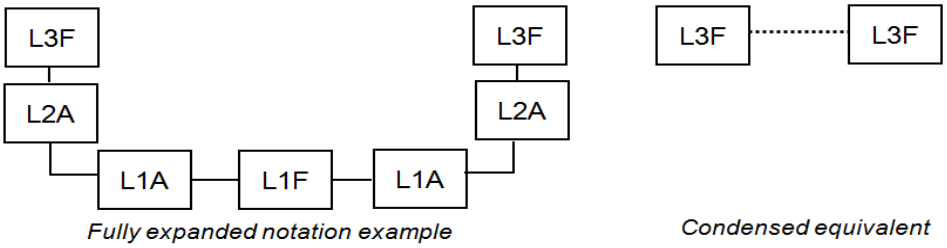

It is also frequently not necessary to draw out all components at a given layer between specific points in a network. This is illustrated with a double line.

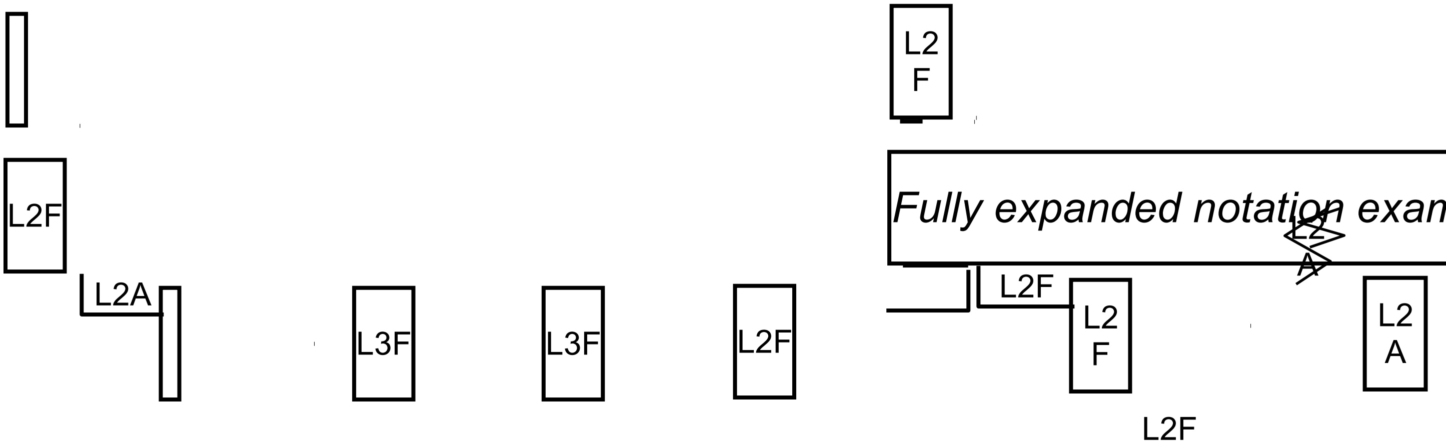

Layer recursion

A consequence of being able to recurse within a given layer and the degree of abstraction in the models is that the normal appearance of hierarchical relationships between the layers may not be obvious. In the diagrammatic example below, L3 is adapted onto L2, mapped onto an L1 link, and then a recursion within layer 2 occurs. An example of this would be an RG mapping IP onto Ethernet over DSL, and at a DSLAM an additional S-tag is added to the Ethernet frame (adapting 802.1Q onto 802.1ad).

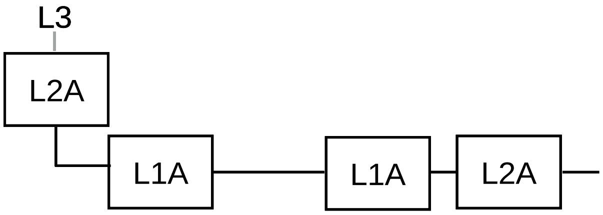

Link Layer Adaptation

Where it is necessary to identify a link layer technology (e.g. DSL) this will be shown as back to back layer 1 Adaptation functions.

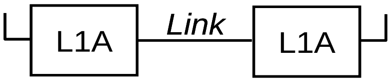

6. Technology Independent Functional Architecture

Having defined the high level reference model with the access and aggregation function sets (Section 4), technology independent functional modules (5.1), and the logical interfaces between the functional modules (4.3.2), this section details the discrete technology independent functional module elements in Aggregation, Access, and Customer Located networks, along with the interfaces between them. Note, not all the functional modules need to be present for a given service or implementation.

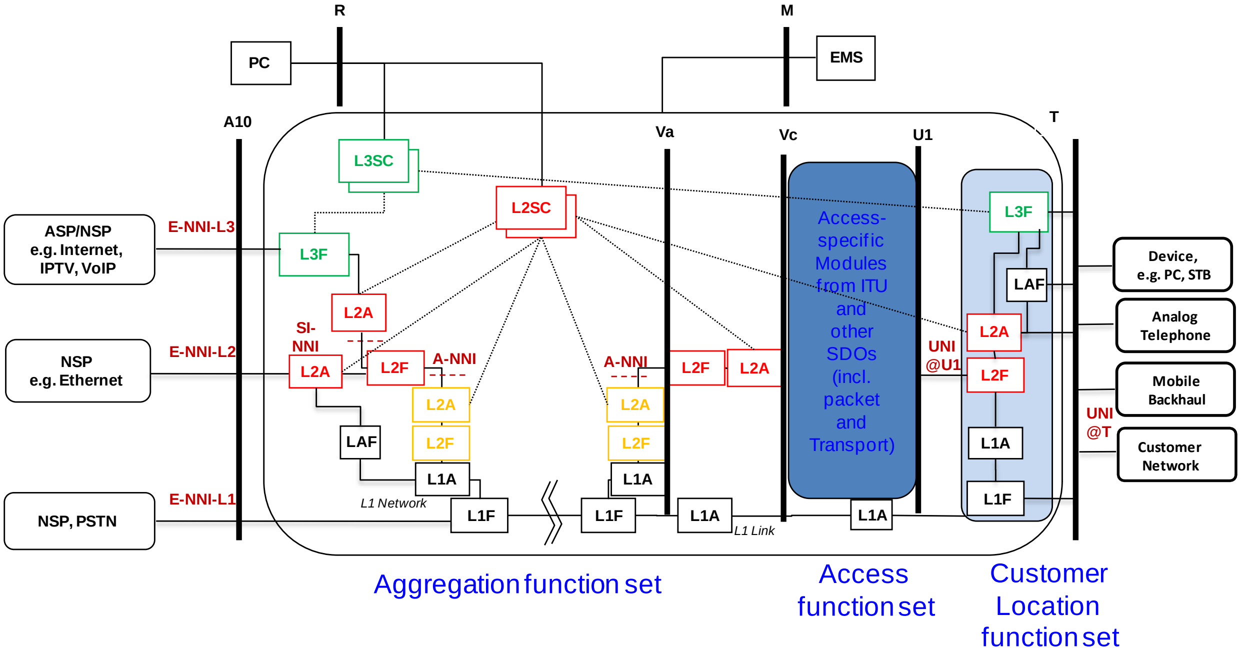

Figure 9 Generic Functional Model

Figure 9 represents a generic Multi-Service Broadband Network functional model using a layering approach. It includes a customer location function set, an access function set, and an aggregation function set, as defined in Section 4.4. The Aggregation function set shows IP aggregation functions, Aggregation functions performed by the Ethernet Service Layer, and optionally aggregation functions performed by the Supporting Aggregation Layer. The Transmission layer is also shown, supporting the higher layers, as well as providing transmission based interconnect. The various logical and business interfaces are shown as well. Dotted lines refer to control plane connections rather than data plane connections in this and other figures in this section. Multiple L2/L3SCs can control several adaptation and forwarding functions, but not all of them are shown for clarity. The STF functions are typically collocated with L2A, L2F and L3F functions but are not shown for clarity.

6.1 Technology Choices for Function Sets and Interfaces

This section describes the aggregation technology options and protocol stacks that need to be supported on the logical and business interfaces.

Figure 10 shows the protocol stacks at the I-NNI interfaces at Va (in case there is a physical Ethernet interface uplink) or V (in case there is an uplink that is performing a supporting aggregation technology).

- Ethernet/IPoE/PPPoE/MPLS
- Circuit Emulation can be performed on the IP and Ethernet Top of Stack technologies

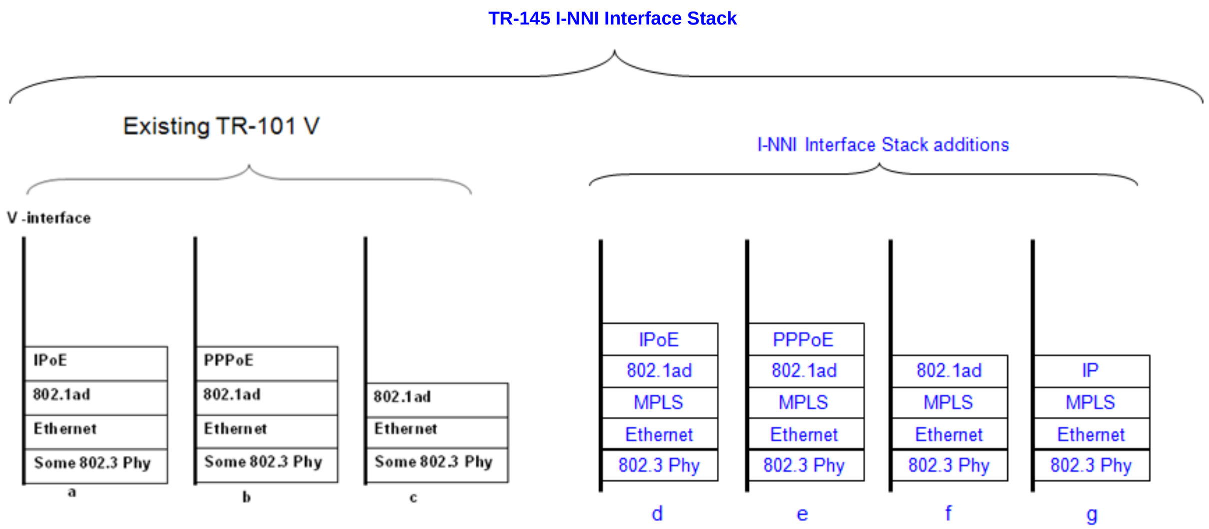

Figure 10 Protocol Stacks at I-NNI interface in the Packet Platform

Figure 11 shows the protocol stacks at the E-NNI interfaces at L1, L2 and L3:

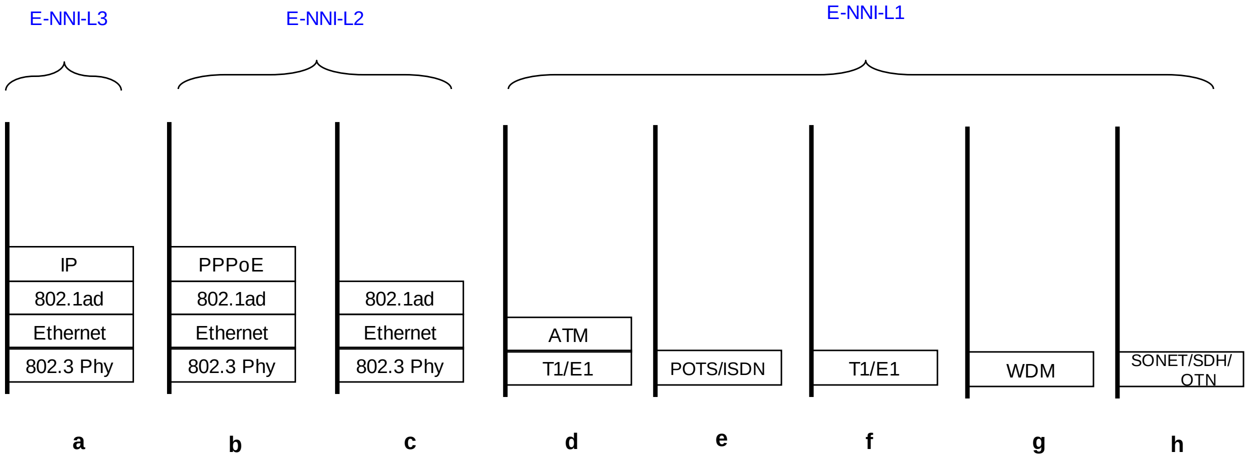

Figure 11 Protocol Stacks at E-NNI interfaces

6.2 Mapping Existing BBF Architectures onto TR-145

This Technical Report provides backward compatibility with TR-101 and TR-156. This section presents how it maps to TR-101 and TR-156 architectures. Other derivatives of the TR-101 architecture, such as TR-167 (GPON fed access node) are supported in the same way.

Figure 12depicts the mapping of TR-101 to TR-145 where the last mile technology is xDSL. As shown in the figure, the TR-101 U DSL interface is distinct from the TR-145 U1 reference point, which is hidden within the B-NT in the TR-101 architecture.

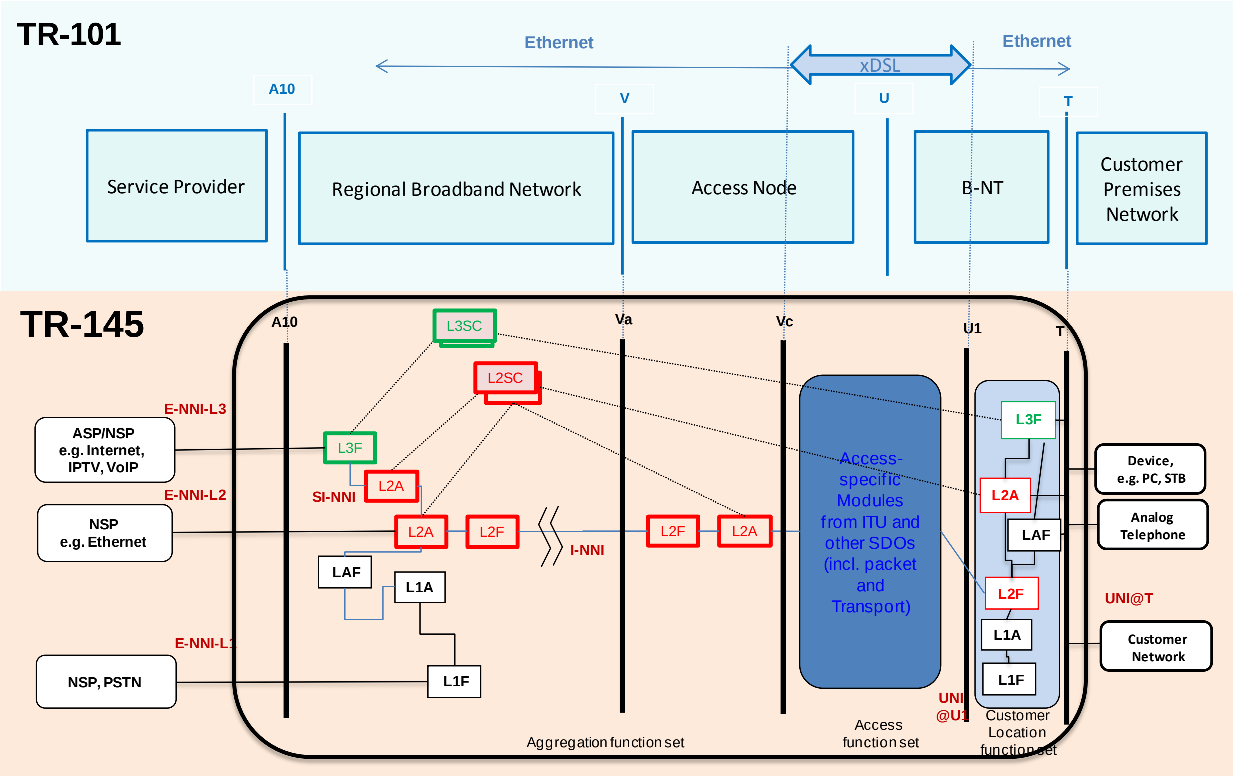

Figure 12 Modelling of TR-101 using TR-145 diagrammatic conventions

TR-101 Access Node requirements and L2 (and above) protocol stacks need to be supported in the architecture between the Vc and Va reference points. Also the TR-101 RG requirements needs to be supported in the architecture by the Customer Location Function Set between T and U1 reference point. Note that xDSL in the picture refers to both DSL ATM and PTM Modes. The dotted lines refer to control-plane connections rather than data plane connections.

Figure 13 depicts the mapping of TR-156 to TR-145 where the ONU hand-off is Ethernet. As shown in the figure, the TR-156 U Ethernet interface is equivalent to TR-145 U1 reference point.

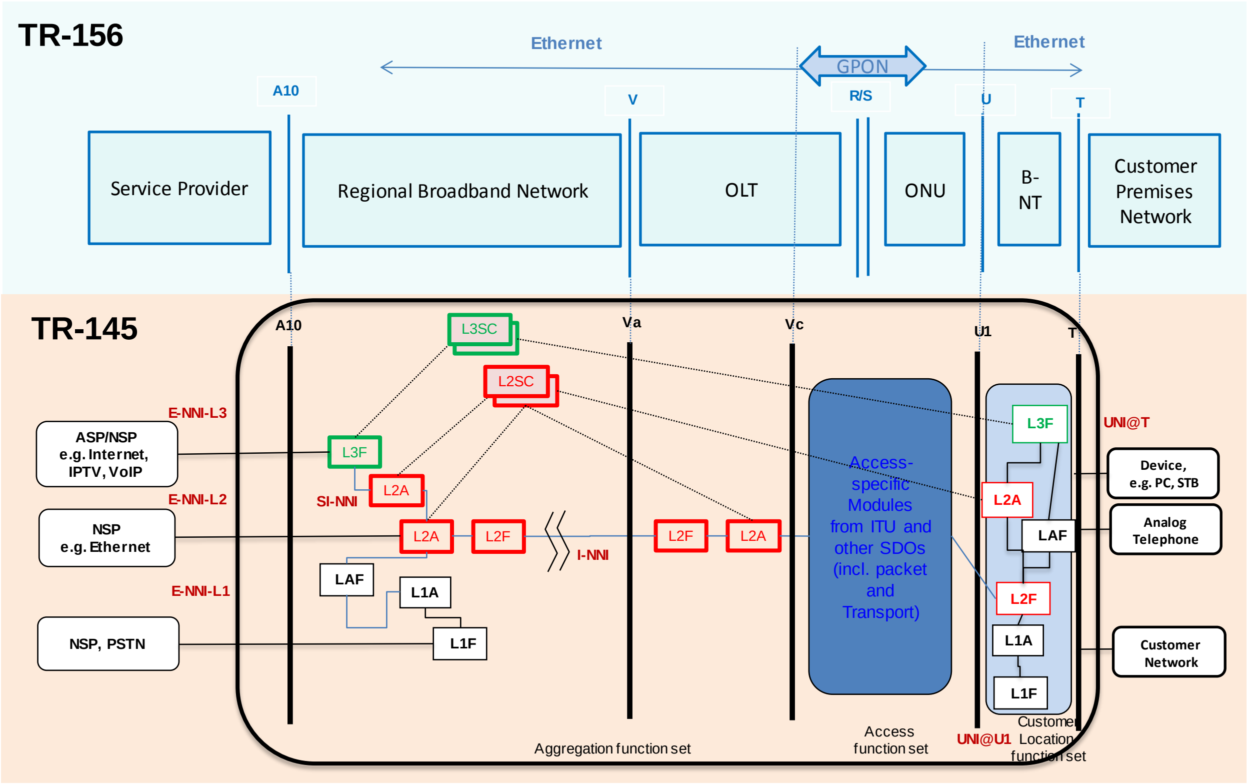

Figure 13 Modelling of TR-156 using TR-145 diagrammatic conventions

Figure 14 depicts an example of the functional model when using ADSL2 as access technology to deliver broadband access services to a consumer and Ethernet service to the business customer.

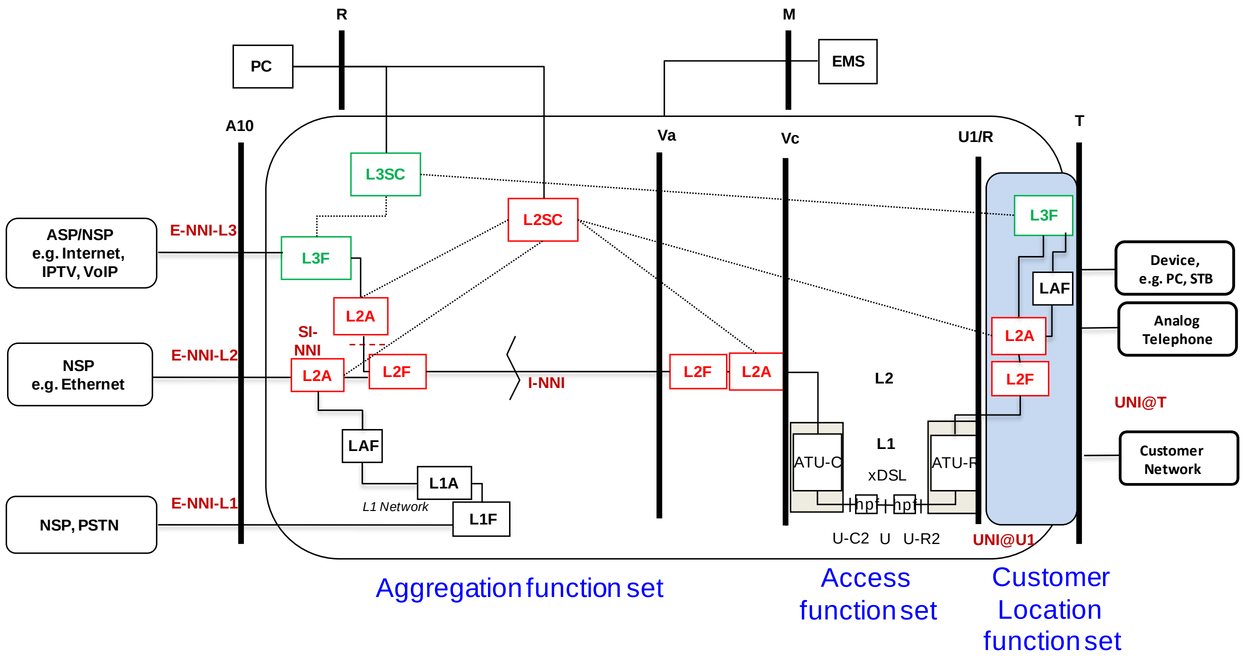

Figure 14 Example of the functional model of TR-101 with ADSL2 as the access technology

6.3 Example considerations for geographical distribution of functionality in an architecture

Multiple Services are supported by the architecture. Therefore most nodes will have to carry traffic for different service types, often received on the same interface. As different Service Types might need a different logical topology in terms of the Ethernet Service Layer as well as the Supporting Aggregation Layer, L2 versus L3 aggregation, Service Edge placement, etc, a node will have to incorporate different functions, on a per service (or set of aggregated or similar services) basis.

A node aggregating customer lines for a plurality of services, will have to be as service agnostic as possible, this is in order to ease provisioning of IVCs, IVC segments, and IVC's terminating into Service Edge functionality. It is also important that a common uplink technology is chosen into the remainder of the aggregation (be it L2 or L3) in order to lessen the need for supporting multiple interface technologies at nodes higher up in the aggregation topology. As the multi-service architecture also needs to support TR-101 access nodes, Ethernet based I-NNI interfaces have to be supported. This does not exclude the use of other virtualization or tunneling techniques (such as MPLS PWs). Hand-off points to NSPs or retail ISPs will be Ethernet or IP over Ethernet except where a legacy TDM/ATM service is supported.

How far the IVC (or a set of IVC segments) reaches into the topology and where it terminates into the L3 aggregation functionality (in other words where the E-NNI is placed) depends on various factors, such as:

- Subscriber density and concentration.
- Whether a high amount of per subscriber/customer control is needed and or the Service Type. In typical scenarios a high amount of per subscriber state calls for greater centralization. Where Qos is only needed on a per service basis this can be more easily distributed (e.g. some Video deployments).
- Failure Domain: distributed Service Edges have a smaller failure domain.
- Ease of provisioning: distributed Service Edges need 'very short' IVCs, thus making the IVC provisioning easier and the stitching of individual segments of the IVCs may be unnecessary. However this places a greater provisioning load on the Service Edges themselves.
- Distribution of Application and Caching Servers (VoD, I-frame, etc), as they need to be placed behind the relevant service edge.
- The opportunity to reduce backhaul oversubscription.
- Providing low latency services might require more distribution of Service Edges.
- Organizational structure of the SP.
- IVC/aggregation technology used: if an IVC technology is used that can leverage an IP based control plane, Service Edge placement is very flexible, as a node can leverage that control plane to built IVCs for service X, as well as being a Service edge for Service Y.
- Regulatory requirements like legal intercept and regulated interconnect points.

In some cases, there might not be a traditional Service Edge within the MSBN e.g. L2VPN business services or L2 wholesale access.

The functional disposition and topology may be service specific and require a variety of technologies, e.g.:

- Point-to-point, point-to-multipoint and multipoint L2VPNs, to build business and wholesale services. For IP Services the same technologies can be deployed to build IVCs interconnecting users to Service Edges.

- Multicast and Unicast IP based aggregation, e.g. to build Video Services.

- L3 VPNs for IP business services and for virtualization of IP services.

A node may need to support more than one of the above technologies at the same time.

6.4 Examples of geographical distribution of functions onto network nodes

6.4.1 Example 1: The 'TR-101 Ethernet Service Layer'

The following example depicts, using the functional modules and interfaces, how a nodal disposition in TR-145 architecture can support the TR-101 Ethernet Service layer. Three nodes are depicted in Figure 15: a Broadband Networking gateway (BNG), an Aggregation Node (AGN), and an Ethernet Access Node (EAN).

At a high level, both the EAN and the AGN perform Ethernet Aggregation functions, whereas the BNG performs both Ethernet and IP Aggregation. Moreover, typical BNGs have MPLS Aggregation functionality as well, although TR-101 does not describe them, as they are not service enabling.

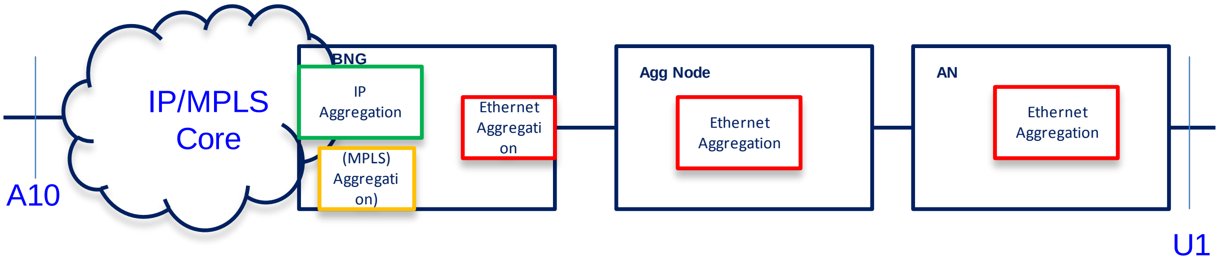

Figure 15 TR-101 High Level Functional Distribution

Figure 16 describes how the various Adaptation and Forwarding functions are placed within the physical nodes, and shows interfaces used to interconnect them.

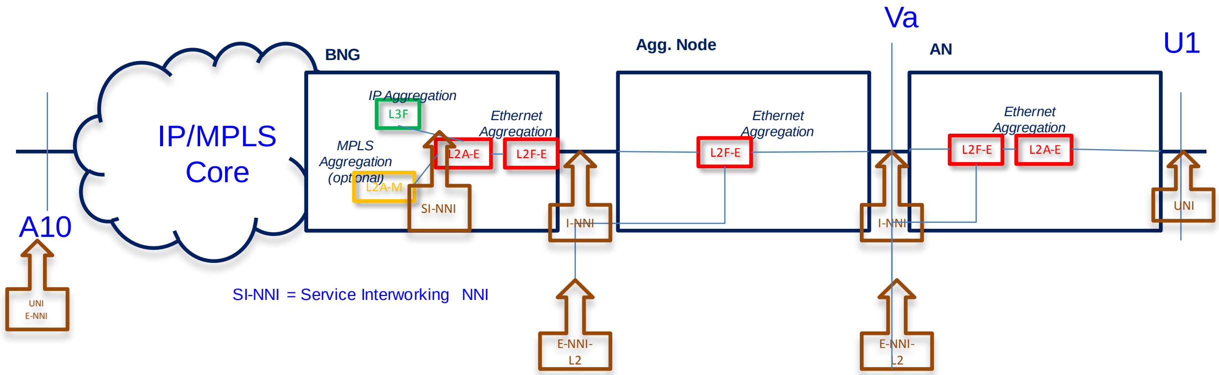

Figure 16 TR-101 Functional Disposition.

As can be seen on the picture, the business interfaces for L2 wholesale can be implemented on the access node, the aggregation node, and also distributed in the network at the A10 reference point. The SI-NNI is an internal logical interface that connects the IP functionality within the BNG with the Ethernet adaptation and forwarding functionality within the BNG.

The TR-101 services are either:

- L2 services between U and A10, or in other words an Infrastructure Virtual Connection between UNI and UNI/E-NNI-L2 at those respective reference points.
- IP Services between U and A10, where the Ethernet Aggregation is performed by an Infrastructure Virtual Connection between the UNI and SI-NNI interfaces.

The requirements for Ethernet Aggregation at the various I-NNIs can be summarized as ‘MUST support 802.1ad forwarding’. The real Service enabling requirements for a TR-101 architecture are at the UNI@U1, the SI-NNI within the BNG, and the E-NNI-L2@A10. Most requirements are for the ‘VLAN manipulation functions’ performed at the U1 reference point i.e. the ability to classify on one incoming VLAN tag (C, or S-tag) and build a S-tag or S/C-tag stack based on that classification. For further details see TR-101i2.

6.4.2 Example 2: Ethernet Access, MPLS Aggregation based architectures

In this example TR-101 access nodes (EAN) can be re-used and connected to MPLS aggregation nodes, which tunnel the Ethernet connections across the network, using VPWS and VPLS techniques. We call this MPLS aggregation node the 'Broadband Aggregation Gateway' or BAG for short. The BNG terminates the MPLS tunnels internally, revealing the Ethernet frames before being subject to IP Aggregation. Figure 17 shows a high level view of the macro sets of functionality needed per node, as well as pointing out where the service enabling requirements will exist.

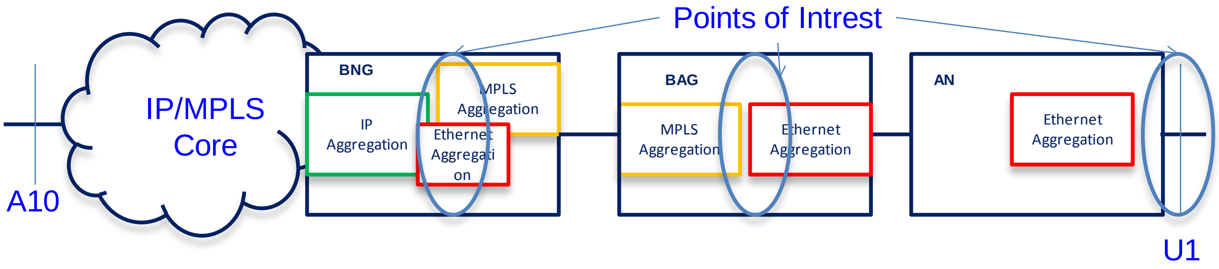

Figure 17 MPLS Aggregation, Ethernet Access, High Level Functional Distribution

Figure 18 describes how the various Adaptation and Forwarding functions are placed within the physical nodes, and shows the interfaces used to interconnect them.

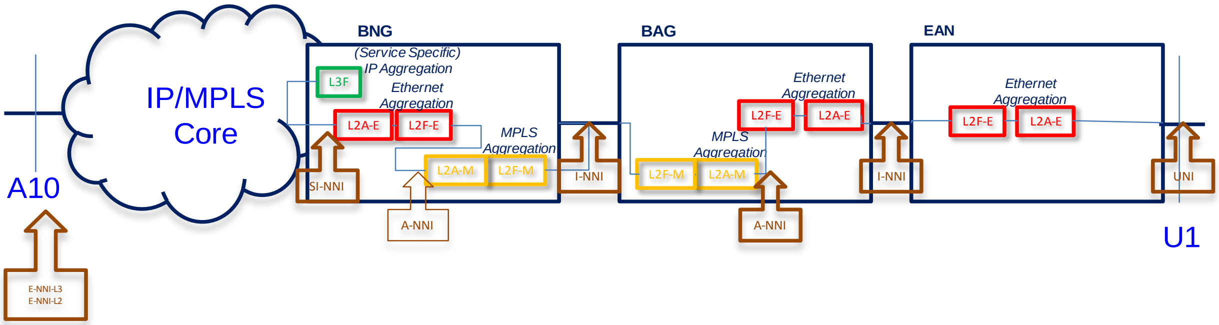

Figure 18 MPLS Aggregation, Ethernet Access, Functional Disposition

Here the two I-NNIs have to support different encapsulation techniques, with the I-NNI between EAN and BAG supporting Ethernet 802.1ad, and the I-NNI between the MPLS aggregation network and the BNG supporting MPLS encapsulation.

The EAN has the same service enabling requirements at the UNI@U1 as the one found in TR-101, as explained in Section 6.4.1 above. The A-NNIs shown on the diagram map 802.1ad tagged frames to MPLS, and can do this based on a classification scheme of up to two tags (S-tag, or S+C-tag). The A-NNI can also support multiplexing various S-tag frames across a single MPLS pseudowire, with optional support for 802.1ad based bridging from the I-NNI into the pseudowire starting at the A-NNI. The terminating A-NNI is positioned inside of the BNG where MPLS de-encapsulation occurs. In this way the SI-NNI interface is exposed to 802.1ad frames for the sake of IP Aggregation.

6.4.3 Example 3: BNG Hierarchies

In this example, TR-101 Ethernet Access Nodes are again re-used. However, the node performing the second level of aggregation can perform both MPLS aggregation as well as IP Aggregation. In other words it is a full-blown BNG and we refer to it as an Edge BNG or BNG$_S$. It will be the IP aggregation point for some services/Service VLANs, whereas it will also perform MPLS tunneling towards more centralized Service Edges. These Service edges we refer to as Service BNGs, or BNG$_{|S}$. They have similar requirements to a BNG/BNG$_E$, but can be specialized with regard to service offerings such as Ethernet Services, Mobility, IPv6, and CGN.

Figure 19 shows a high level view of the macro sets of functionality needed per node, as well as pointing out at which interface the service enabling requirements will be put forward.

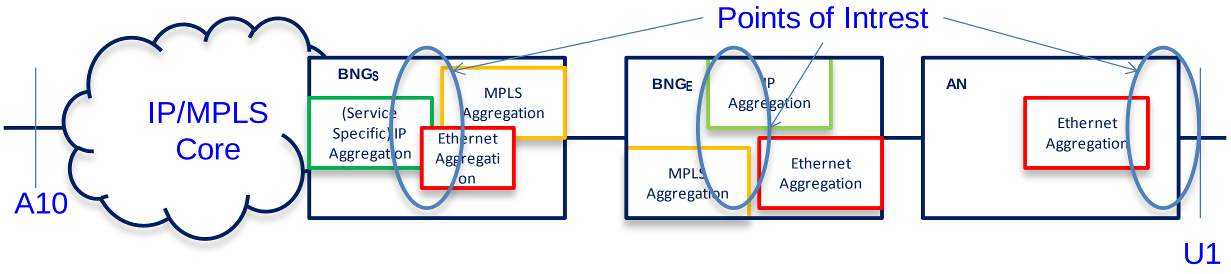

Figure 19 BNG Hierarchies, High Level Functional Distribution

Figure 20 describes how the various Adaptation and Forwarding functions are placed within the physical nodes, and shows interfaces used to interconnect them.

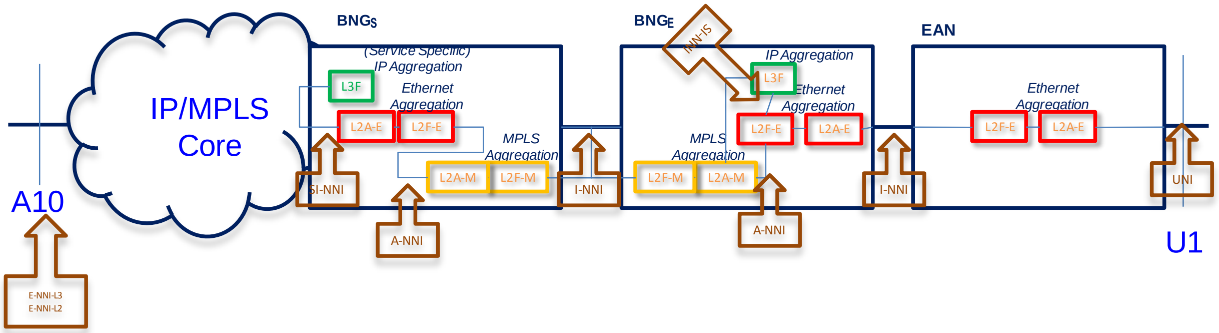

Figure 20 BNG Hierarchies, Functional Disposition

Here the two I-NNIs have to support different encapsulation techniques, with the I-NNI between EAN and BAG supporting Ethernet 802.1ad, and the I-NNI between the MPLS aggregation network and the BNG supporting MPLS encapsulation.

The EAN has the same service enabling requirements at the UNI@U1 as the one found in TR-101, as explained in Section 6.4.1.

The BNG$_E$ performs Ethernet Aggregation and can either forward packets via MPLS pseudowires, or through IP Aggregation/routing. The de-multiplexing can be done on a per VLAN basis, or on a per destination MAC-address basis through 802.1ad bridging within the L2F-E. The A-NNI interface takes care of the MPLS tunneling and has the same requirements as the A-NNI of the BAG as per Section 6.4.2. The SI-NNI has similar requirements than the BNG as per Section 6.4.1.

The BNG$_S$ performs similar MPLS de-encapsulation, Ethernet and IP Aggregation functionality as a BNG$_E$, but does not need to switch through all core-facing MPLS traffic.

6.4.4 Example 4: MPLS based Aggregation in the Access Nodes

In this example the AN implements Ethernet Aggregation and MPLS aggregation function while performing at a minimum the TR-101 Ethernet Service Layer. For the sake of differentiating it from the EAN we refer to this node as the MPLS AN or MAN. The upstream nodes have to perform MPLS forwarding, a subset of the functionality of the BAG, while supporting extensions needed to scale the MPLS dataplane and control plane requirements on the Access Nodes downstream. The BNG has similar Ethernet, MPLS and IP Aggregation functionality as described in Example 2.

Figure 21 shows a high level view of the macro sets of functionality needed per node, as well as pointing out at which points the service enabling requirements will be put forward.

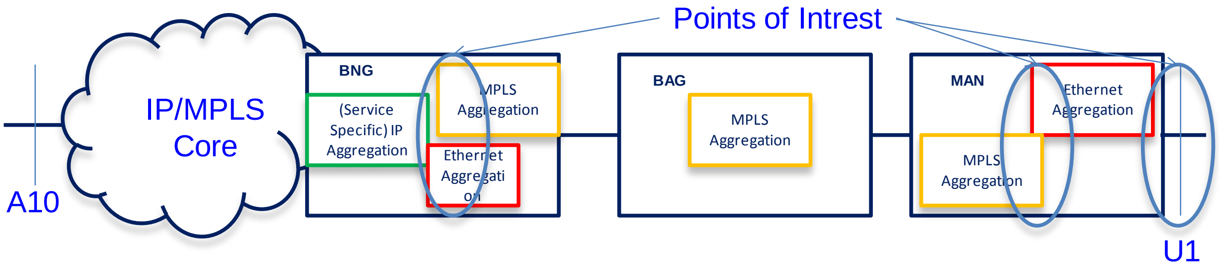

Figure 21 MPLS to the Access, High Level Functional Distribution

Figure 22 describes how the various Adaptation and Forwarding functions are placed within the physical nodes, and shows interfaces used to interconnect them.

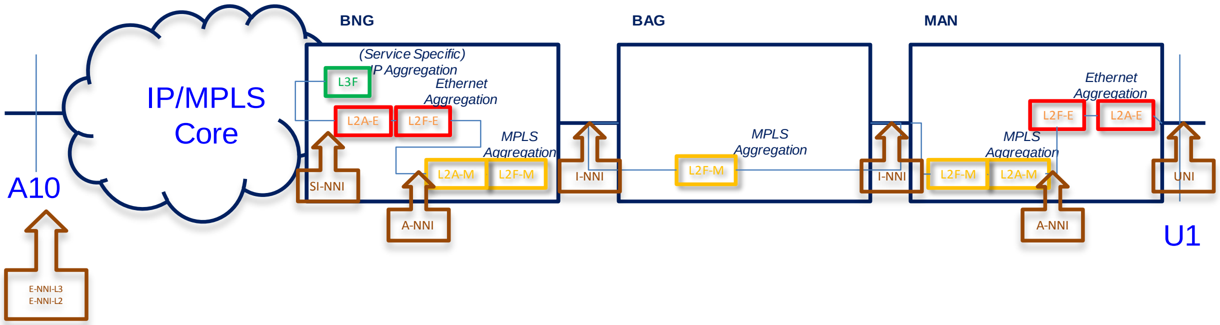

Figure 22 MPLS to the Access, Functional Disposition

Here the A-NNI performs the same requirements as the A-NNI on the BAG in Section 6.4.2, except that it is now performed within the Access Node. The BNG has the same requirements on the I-NNI, SI-NNI and A-NNI as in Section 6.4.3. The BAG only has to perform a subset of the requirements laid out in Section 6.4.2 on the I-NNI while supporting extensions needed to scale the MPLS dataplane and control plane requirements on the Access Nodes downstream.

6.4.5 Example 5: MPLS and IP based Aggregation in the Access Nodes

In this example, the Access Node has IP Aggregation, MPLS Aggregation and Ethernet Aggregation functionality. Essentially the TR-101 architecture is collapsed inside one node-type. We refer to this node as the BNG embedded Access Node or BAN. All nodes upstream are performing a subset of the functionality of the BAG out of in Section 6.4.2, while supporting extensions needed to scale the MPLS dataplane and control plane requirements on the Access Node. Optionally downstream BNG/BNG$_S$'s can be deployed for services not terminated at the access node.

Figure 23 shows a high level view of needed macro sets of functionality per node, as well as pointing out at which point the service enabling requirements will be put forward.

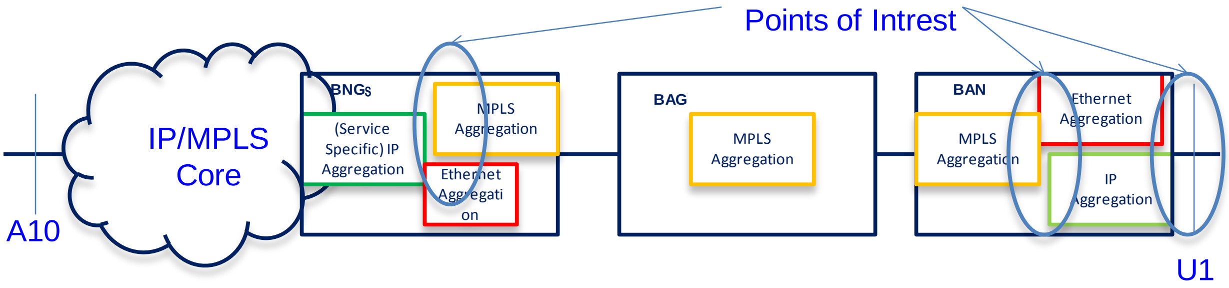

Figure 23 BNG in the Access, High Level Functional Distribution

Figure 24 describes how the various Adaptation and Forwarding functions are placed within the physical nodes, and shows interfaces used to interconnect them.

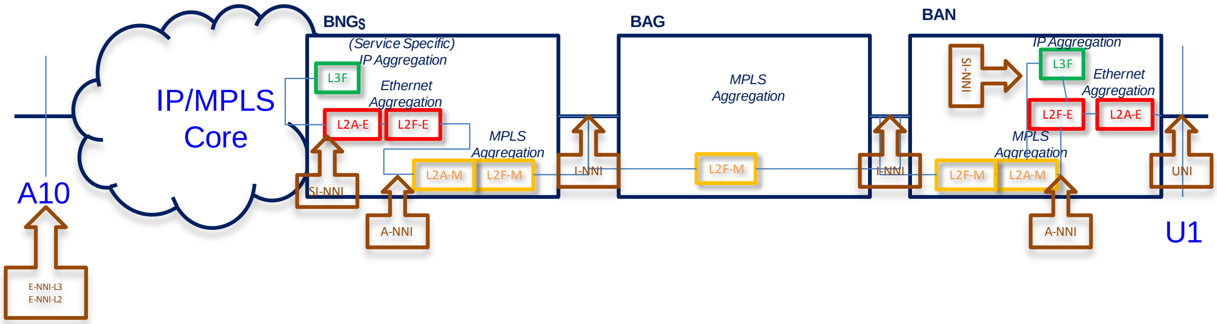

Figure 24 BNG in the Access, Functional Disposition

As far as functionality is concerned, the requirements at the A-NNI and SI-NNI interfaces are the same as in the previous examples, only the nodal disposition has changed, as they now are all within the BAN. Requirements at the I-NNI are similar to the architecture described in Section 6.4.4.

6.5 Conclusion: Towards WT-178

As can be seen in the previous examples, the functional modules and interface descriptions can be used to create multiple architectural options for future work i.e. WT-178.

A lot of the requirements at the various interfaces (UNI, I-NNI, E-NNI, SI-NNI and A-NNI) can be re-used across nodes, and even across architectural options. Macro sets of requirements can be grouped together in order to aid this re-useability.

7. Requirements

7.1 Requirements for Service Layers

7.1.1 Support of lower layers

Lower layer OAM signals need to be converted to meaningful events/alarms to the upper layer adaptation functions, such that they can react appropriately when a lower level event occurs. An example could be where a G.709 L1A adaptation function, upon receiving a G.709 OAM Loss of Frame event, needs to inform its collocated L2A function to generate a AIS to its peer L2A across the physical link.

R-1 The L1A function MUST support generating the appropriate signals at the L2A functions and associated L2A peers in case of L1 Loss of Continuity.

R-2 The TR-145 Architecture needs to support OAM interactions between Layers

7.2 The Ethernet Service Layer

The functional modules can have different incarnations depending on the place in the network (geographical distribution, whether they perform in the Ethernet Service Layer or the Supporting Aggregation Layer and depending on the ingress and egress logical interfaces it uses to interconnect with other modules. One example is the L2A function that lives within a TR-101 access node: here the ingress logical interface is the UNI@U, which is defined as an Ethernet-over-xDSL interface including both ATM and PTM (Packet Transfer Mode) framing, and the egress interface is the I-NNI@Va, which is defined as an Ethernet/802.1ad interface. Most of the requirements for this L2A incarnation can be found in TR-101 and are not repeated in this document. This document only addresses requirements beyond TR-101.

Similarly the TR-101 BNG Ethernet Functionality as well as the IP aggregation functionality (L3F @ SI-NNI interface) will have to be supported through TR-145 functions and interface descriptions, more specifically the L2A between I-NNI and SI-NNI (Ethernet functions) and the L3F between the SI-NNI interface and E-NNI-L3 (IP functions).

Essentially TR-101 had two types of services: L2VPN/TLS Services and IP Services. An Ethernet Aggregation Layer that leveraged 802.1ad was used to built L2VPN/TLS type services, as well as providing connectivity for IP Services between RGs and BNGs.

See Section 6.4.1 for an example of how TR-145 functionality and interfaces can be used to build the Ethernet Service Layer for TR-101.

All TR-145 compatible architectures have to be able to deliver this ‘Ethernet Service Layer’, independent of any underlying and supporting aggregation and tunnelling technologies. These supporting functionalities are also done through L2A functionalities, but this time across the A-NNI interface.

The architecture also needs to support adaptation of legacy services (TDM) to the supporting L2F functionality

The Ethernet Service Layer also needs to support IPv6, as described in TR-177 and TR-187 and IPv6 multicast as described in Section 5.2.

It is possible to create general requirements for the Ethernet Service Layer, which are essentially IVCs between the UNIx and E-NNI-L2 interfaces, or IVCs between the UNIx and the SI-NNI interface that is facing the IP Aggregation functionalities

R-3 The Ethernet Service Layer IVCs (between UNI@U1 and E-NNI-L2@A10 or between UNI@U1 and SI-NNI) MUST be able to support the Services and Requirements of TR-101.

R-4 The Ethernet Service Layer IVCs (between UNI@U1 and E-NNI-L2@A10 or between UNI@U1 and SI-NNI) MUST be able to support the Services and Requirements of TR-177 and TR-187.

R-5 The Ethernet Service Layer IVCs (between UNI@U1 and E-NNI-L2@A10 or between UNI@U1 and SI-NNI) MUST be able to support the services described in TR-144 [4].

R-6 The Ethernet Service Layer and supporting technologies MUST be able to support IPv6 multicast.

R-7 The Ethernet Service Layer IVCs (between UNI@U1 and E-NNI-L2@A10) MUST be able to support emulated TDM and ATM services through the use of circuit Emulation over Packet techniques.

MEF 20 describes the requirements for an Ethernet business service UNI Type 2.2. The architecture needs to support these on the UNI interfaces at the U and/or T reference points. Functionality includes service instance mapping, Ethernet LMI, traffic profiles, and security.

R-8 The Ethernet Service Layer IVCs (between UNI@U1 and E-NNI-L2@A10 or between UNI@U1 and SI-NNI) MUST support the MEF 20 UNI Type 2.2 as described in [10]

R-9 The L2A Functions MUST provide OAM mechanisms for path, node and link failure detection between various interfaces.

R-10 The L2A Function MUST be able to redirect upstream flows when a switchover occurs between the active and standby L3F function

R-11 The IVCs MUST be able to interact with Policy Control in order to set up multicast configuration The STF between Va and Vc MUST be able to interact with RC (Resource Control) to admit/reject the multicast channel on a per UNI basis and update the resource pool in the RC for multicast traffic in the Ethernet Service Layer.

Sometimes technologies such as Ethernet over SONET/SDH are leveraged to extend the reach of user facing or network facing interfaces. However, there is often a speed mismatch. Policing or making use of techniques such as Ethernet Pause (802.3x) can often lead to very choppy traffic

traffic-patterns. Shaping e.g. a 1Gbps Ethernet to 622 Mbps is a preferred solution as there will never be a need for the Ethernet over SONET equipment to police/pause the traffic.

R-12 The STF function MUST be able to shape traffic to the available bandwidth of the underlying physical path (which may span one or more hops) in order to achieve the appropriate QoS/CoS goals.

TR-101 does define functionality for adding S-VLANs depending on the received C-VLAN but this requirement is split in two for TR-145: One requirement that describes 'Selective' double-tagging (i.e. when receiving a given C-VLAN, or a given set of C-VLANs, or a given range of C-VLAN, or a given set of ranges of VLANs, add a given S-VLAN). The other requirements describes 1-to-2 translation: translating a single tagged frame with C-VLAN x , to a double-tagged frame with S-VLAN w and C-VLAN y. However, the following two requirements are written in a technology unspecific manner i.e. the tag added is denoted as a 'service instance tag' which could be an S-VLAN, an I-tag, a set of MPLS labels, etc.

R-13 The L2A function at the UNI@U1 and I-NNI and E-NNI (at Va) interfaces MUST support selective service instance tagging in the upstream direction. A configurable service instance tag MUST be added when receiving a given C-VLAN, or a given set of C-VLANs, or a given range of C-VLAN, or a given set of ranges of C-VLANs. This operation MUST operate in a symmetric fashion for downstream and upstream operation.

R-14 The L2A function at the UNI@U1 and I-NNI@Va and E-NNI@Va interfaces MUST support 1-to-2 service instance tag translation in the upstream direction. A configurable service instance tag MUST be added AND a configurable C-VLAN MUST replace the initial C-VLAN. This operation MUST operate in a symmetric fashion for downstream and upstream operation.

OAM requirements based on 802.1ag CFM are already captured in TR-101 and will not be repeated here. However, performance management leveraging Y.1731 is not captured in TR-101. Y.1731 builds upon leveraged 802.1ag / CFM as a transport for performance monitoring and management.

R-15 The L2A function in the Ethernet Service Layer MUST support the appropriate aspects of Performance Management as per Y.1731

7.3 Requirements for technology independent functional modules in the Supporting Aggregation Layer

This section captures the requirements for technology independent functional modules in the Supporting Aggregation Layer. It is observed that, in order to provide useful and clear requirements for a functional module, the requirement always has to be stated referencing a logical interface connecting the module.

7.3.1 Requirements for Ethernet Flow Points on A-NNIs connecting to an L2A

7.3.1.1 VLAN Classification Requirements

This section focuses on the L2A functionalities residing at both sides of the A-NNI interface, as seen in Figure 25. The L2A-E functionality supports the Ethernet Service Layer by implementing Ethernet while the L2A-X functionality adapts Ethernet to the supporting tunnelling technology of choice. L2A-X and L2F-X refer to the functionality in the Supporting Aggregation Layer.

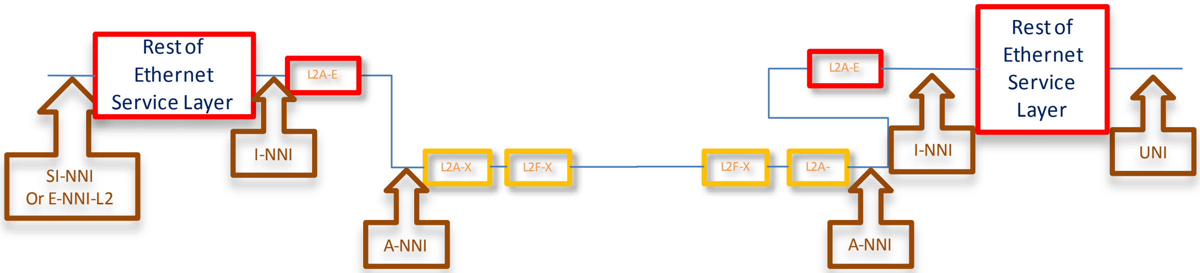

Figure 25 Ethernet Service Layer

As seen on the diagram, the I-NNIs in the Ethernet Service Layer only have to support 802.ad framing, while the A-NNI interface has the adaptation features needed to adapt the Ethernet Service Layer onto a supporting layer (like MPLS).

Note that 'Rest of Ethernet Service Layer' denotes any Ethernet Service Layer specific L2A functionalities upstream or downstream from the tunnel between the two A-NNIs. For example various L2A's can be instantiated in a ring or daisy-chained based access network.

This leads to the following requirements:

R-16 An I-NNI interface interconnecting two L2A functional modules MUST support Ethernet 802.1ad based VLAN tagging

Section 4.5 describes the concept of an IVC and an example IVC that creates an association between two xNIs. Multiple IVCs can terminate on an SI-NNI, UNI or E-NNI. A single instantiation of an IVC is referred to as an Ethernet Flow Point (EFP). This section defines requirements on a per EFP basis that can drive functionality needed inside this incarnation of the L2A. The below requirements for EFPs apply on A-NNIs.

Here EFPs are classified as a L2 traffic flow containing a combination of S and C-VIDs. Two EFPs on two different physical incarnations of the A-NNI that have the same classification criteria are in fact two different EFPs for the L2A functional module and should be treated as thus. This allows per port VLAN local significance.

R-17 EFP classification MUST support untagged, priority tagged, single-tagged (802.1ad S-VLAN) and double-tagged frames(802.1ad S/C VLANs)

R-18 Different EFPs, all with different classifications MUST be able to co-exist on an A-NNI interface interconnecting two L2A functional modules.

R-19 EFPs with the same VLAN tagging classification, but which are on different physical ports MUST be treated as independent EFPs.

EFP classification needs to be flexible enough to support unique VLANs, ranges of VLANs, lists of VLANs, lists of ranges of VLANs, as well as a wildcard, and this for up to 2 VLAN tags inside the packet, as well as untagged traffic. The classification is non-exact, meaning it will not look beyond the VLAN tagging criteria.

R-20 If an EFP matches S-VLAN tagged traffic, the classification criteria MUST support a unique S-VID, a range of contiguous S-VIDs, a list of S- VIDs, or a list of ranges of S-VIDs, disregarding the inner tag should it exist.

R-21 An EFP matching untagged traffic MUST be supported.

R-22 An EFP matching a unique S-VLAN, the classification criteria for matching the inner C-VID MUST support matching a unique inner C-VID, a range of inner C-VIDs, or a list of inner C-VIDs

Different EFPs, all with different VLAN classification can coexist on the A-NNI interface interconnecting two L2A functional modules. Traffic classification should perform a longest match lookup through the list of configured EFPs i.e. if two EFPs match the same traffic flow with VID X, but one only matches a range of VLANs that includes VID X, where the other matches VID X exactly, packets are classified as belonging to the second EFP. Also traffic classification should perform a non-exact match i.e. if a packet is double-tagged, and an EFP only classifies on the outer (or only) tag, it should still be possible to classify this packet into this particular EFP.

R-23 If EFP classification overlaps, a method MUST be provided to classify a given traffic flow into a unique EFP.

It is useful to configure the behaviour when a given traffic flow is not matched by any EFP on the A-NNI interface interconnecting two L2A functional modules. The default behaviour should be to drop that traffic. Alternatively a catch-all EFP should be configurable.

R-24 Traffic not matching any EFPs on an A-NNI MUST be able to be dropped

R-25 It MUST be possible to configure a ‘default EFP’ for all traffic that has not been matched by another EFP.

An alternative way to classify traffic into an EFP is to use Class of Service. The Use cases would be to allow traffic that is received on one VLAN to get assigned to different EFPs based on CoS. Another useful alternative is classification based on EtherType. One such Use case would be to separate PPPoE from IP traffic without worrying about the received VID.

R-26 EFP classification SHOULD support classifying traffic into different EFPs based on the CoS.

R-27 It MUST be possible to support EFP classification on the basis of the following received EtherTypes
* IPv4/ARP (0x800,0x806),
* PPPoE (0x8863,0x8864)
* IPv6 (0x86DD)

VIDs can be manipulated through translation, or pushing a VID, or popping a VID, (pushing a VID allows selective double-tagging behavior), taking into account the need for symmetry in the VLAN manipulation actions.

R-28 After VID based classification, VID translation MUST be supported per EFP The following translations MUST be supported:

* S-VID/C-VID to 1 S-VID
* S-VID to S-VID/C-VID
* S-VID to S-VID
* S-VID/C-VID to S-VID/C-VID

Note: VLAN translation is only appropriate on unique S-VIDs or C-VIDs

R-29 After VID based classification, VID popping pop MUST be supported. Up to two VIDs (S-VID or S/C-VID) can be popped (or removed). VID removal MUST NOT depend on the configuration of a given VID. VLAN popping MUST be only configurable on unique S-VIDs or C-VIDs i.e. not on ranges or lists.

(Note: before removing/popping a VID, R-41 below might apply.

R-30 After VID based classification, VID pushing MUST be supported, by explicitly configuring the VID to be pushed onto the packet. VLAN push MUST only be configurable on unique S-VIDs or C-VIDs, i.e. not on ranges or lists.

(Note: before removing/popping a VID, R-41 below might apply.

R-31 Both VID translation as well as VID pop and push operations MUST, by default, be symmetrical with respect to the traffic received on and transmitted out of an EFP on an A-NNI interface interconnecting two L2A functional modules.VID translation SHOULD be configurable to allow for asymmetric mapping on a per EFP basis.

After classification and VLAN manipulation EFPs are assigned or multiplexed into a L2F specific service instance. This can be on a one to one basis (one EFP is assigned to one L2F specific service instance), or on a many to one basis (multiple EFPs are multiplexed onto a single L2F specific service instance). The latter can be done by retaining or adding an IVC-wide identifier e.g. a VLAN id, or by bridging (learning source MAC addresses and forwarding to destination MAC addresses). The L2F function can be as simple as a point to point connection between EFPs.

R-32 EFPs on an A-NNI interface interconnecting two L2A functional modules MUST be assignable to a unique L2F specific service instance e.g. to map into a tunnel. Note: this can be done by retaining or translating to a unique S-VID per EFP before handing over to the L2F specific service instance.

R-33 Multiplexing EFPs on a single unique L2F specific service instance MUST be supported by using bridging (i.e. MAC-address based forwarding) as the multiplexing technique (assign traffic from a L2F specific service instance to an EFP based on destination MAC -address, learn source MAC -addresses from traffic received on an EFP, flood if the destination MAC -address is unknown).

R-34 Interconnecting two EFPs across an L2A MUST be supported. This allows to interconnect two EFPs without requiring an L2F (bridging).

If bridging is deployed, multicast and isolation between EFPs has to be taken into account. Also bridging allows a portion of the received traffic to get forwarded to the L2F specific service instance, but also to send traffic towards an L3F functional module, again based on the MAC-address tied to the L3F instance.

R-35 If Bridging between EFPs and the L2F specific service instance is configured, the bridging instance MUST be able to support constrained multicast forwarding through the use of IGMPv3 snooping.

R-36 To support Pv6 across the Supporting Aggregation Layer, an A-NNI interface interconnecting two L2A functional modules MUST support the identification and processing of MLD messages.

R-37 If Bridging between EFPs and the L2F specific service instance is deployed as part of the functionality of the A-NNI interface interconnecting two L2A functional modules, split horizon forwarding MUST be supported across the EFPs

The following are requirements resulting from the Use Case 0

R-38 Configuring an EFP at the A-NNI interface SHOULD allow configuring functionality for an associated remote EFP using the appropriate L2SC functionality.

7.3.1.2 Requirements for traffic filtering and scheduling (STF) on EFPs at A-NNI and E-NNI interfaces

This section defines requirements for the Scheduling and Traffic Filtering functions, on a per EFP basis. It is observed that, at least for these requirements an EFP has an L2A and an STF part. Both of these functions are part of the forwarding plane. An STF function is therefore always collocated with an L2A function. Both of these functions can be set up by manual configuration, or may utilize a L2SC function that can glean data plane and control plane packets, interact with a policy system, and reconfigure the L2A and STF functions accordingly.

Based on the previous discussion it can be seen that EFPs are a kind of 'L2 sub-interface' of the A-NNI/NI-NNI logical interface. Any traffic filtering or scheduling can therefore be done on a per EFP basis. In the remainder of the document, every EFP has an instance of the Scheduling and Traffic Filtering functional module (STF). This STF can support manual configuration, or can be auto-provisioned through a session control functional module (L2SC), that uses the data and control plane traffic received on those EFPs as a trigger. The L2SC functional module will interrogate the policy control functional module (PC) before configuring the STF functions associated with each EFP.

R-39 The STF function MUST be manually configurable i.e. without requiring interaction with the L2SC functional module.

R-40 The STF function SHOULD be configurable through the L2SC Functional Module.

As Ethernet packets (IPoE, PPPoE, IPv6oE) are received, the priority fields in those packets have to be mapped onto the L2A/F specific technology. This can be done by looking at the Ethernet priority fields, looking at higher layer information such as IP DSCP, or by doing a per EFP classification

R-41 The STF function MUST support mapping the received per EFP 802.1p priority values (within the S-tag or S/C-tag) into the L2A/F specific priority field. This MUST be done before any pop or push operation. The reverse operation MUST also be supported.

R-42 The STF function SHOULD support mapping the received per EFP EtherType values into the L2A/F specific priority field. Note: Ethertype refers to the Ethertype behind the S-Tag or S-tag/C-tag, should they exist.

R-43 The STF function MUST support setting the L2A/F specific priority field on a per EFP basis i.e. without looking at the packet content

R-44 The STF function MUST support mapping the received per EFP IP DSCP values into the L2A/F specific priority field.

R-45 The STF function MUST support a 1-rate-3-color (RFC2697) Policer per EFP and CoS within the EFP

R-46 The STF function MUST support 2-rate-3-color (RFC2698) Policer per EFP and CoS within the EFP

The following are the QoS requirements, on a per EFP basis, for the A-NNI interface interconnecting two L2A functional modules.

R-47 The STF function at the A-NNI and the E-NNI interfaces MUST support applying H-QoS.

R-48 Each H-QOS instance in R-47 above MUST be able to perform per EFP shaping

R-49 Each H-QoS instance in R-47 above MUST support traffic classification.

R-50 Within the context of the H-QoS instance attached to an EFP, traffic in that EFP belonging to a certain class MUST be assignable to a queue with a minimum bandwidth guarantee. Note: the minimum can be zero

R-51 Within the context of the H-QoS instance attached to an EFP, traffic in that EFP belonging to a certain class MUST be assignable to a queue. These queues can be configured with a given forwarding behaviour such as Expedited Forwarding, Assured Forwarding, etc.

R-52 Per Class Queuing of traffic MUST be supported within the per EFP shaped instance.

R-53 Scheduling in classes within an EFP MUST be work-conserving i.e. the scheduling tries to avoid link resources to go unused

R-54 It MUST be possible to shape a configured set of EFPs

Traffic filtering: This includes limiting the number of MAC -addresses per EFP, or per bridging instance if EFPs are multiplexed using bridging. MAC and IP ACLs need to be supported while forwarding at L2. DHCP snooping with automatic IP and ARP filtering needs to be supported as well.

R-55 The STF function MUST support limiting the number of MAC-addresses per EFP if bridging between EFPs is used.

R-56 The STF function MUST support limiting the number of MAC -addresses per L2A/L2F instance if bridging between EFPs is used. Once the limit is reached, no more MAC -addresses are learned and associated with EFPs.

R-57 The STF function MUST support a MAC ACL (Source/destination/Ethertype) per EFP if bridging between EFPs is used.

R-58 The STF function SHOULD support a L3/L4 ACL per EFP (IP protocol numbers and UDP/TCP port numbers).

The following requirements are all within the context of the deployment of bridging between EFPs

R-59 The STF function, co-located with each EFP, MUST support DHCP filtering per EFP in order to prevent DHCP server originated messages being received from the network on an EFP.

R-60 The STF function, co-located with each EFP, MUST support DHCP snooping in order to install an IP Address to MAC Address binding table.

R-61 The STF function, co-located with each EFP, MUST support dynamically installing IP filters based on the DHCP snooped information in R-64 in order to prevent upstream traffic from IP hosts that have not been identified through DHCP snooping.

R-62 The STF function, co-located with each EFP, MUST support filtering ARP packets based on the DHCP snooped information in the previous requirement. It MUST drop any unicast ARP packet that does not have the proper IP to MAC binding. It MUST ensure that a downstream broadcast ARP packet is sent out only on the appropriate EFP.

Multicast Requirements

R-63 The STF function MUST be able to filter unwanted multicast traffic:

- It SHOULD be possible to control at per-interface level if multicast traffic will be forwarded from this interface and which multicast groups, (S,G) or (*,G), are allowed.
- it SHOULD be possible to control at per-interface level if multicast traffic will be forwarded to this interface and which multicast groups, (S,G) or (*,G), are allowed.
- The STF function MUST interact with **RC** to update the resource pool in the **RC** for multicast traffic in the IP Aggregation Functional Set

7.4 Customer Location Functions (CLF)

Customer Location Functions have been modelled in this document but actual requirements are out of scope of this document. Examples of Functionality needed in the CLF can be found in the below list, but this list is not a definite list.

- Functionality out of existing BBF TRs:
  - L3F: router, DHCP, NAT functions, as defined in TR-124i3
  - L2A: Address resolution (ARP), VLAN tagging
  - L2F: bridging and forwarding
  - L3SC: PPP session, IP Sessions …

- New Functionality:
  - L2A: VLAN architecture and mapping
  - L2SC: 802.1X, SLA, performance and monitoring, and OAM
  - L1F: wavelength, direct fiber, and direct copper access
  - LAF: legacy Adaptation, e.g. ATA for voice and CES and tunneling for low speed TDM or legacy ATM, etc.
  - Synchronization

The CLF needs to be able to support IPv4 and IPv6 multicast.

7.5 Requirements for L2SC

The following are generic requirements that apply to L2SC at both the Ethernet Service Layer as well as the Supporting Aggregation Layer, unless it is noted otherwise.

R-64 L2SC MUST support redundant path routing between redundant IVCs terminated in widely separated geographical locations.

R-65 L2SC MUST support simple redundant IVC switchover without geo diversity

R-66 L2SC MUST support a dynamic Ethernet session establishment mechanism that enables end to end session (L2 and L3) establishment between endpoints and multiple service edges.
Note: Ethernet session established by L2SC can be an end to end or a segment IVC.

R-67 For the case of end to end L3 IP session, L2SC MUST establish an Ethernet session prior to establishment of corresponding L3 sessions.

R-68 L2SC MUST support multiple Ethernet sessions over the same UNI interface at U1 or T

R-69 L2SC MUST be able to recognize a user Ethernet endpoint attachment request, e.g. IEEE 802.1x request message at UNI interface at U1 or T

R-70 L2SC MUST be able to issue a AAA request to PC to authenticate the attachment.

R-71 L2SC MUST be able to issue an authentication request to PC to authenticate an Ethernet segment session per the request from a different Service Layer, e.g. Supporting Aggregation Layer.

R-72 L2SC MUST be able to receive and act upon an accept/reject response authentication from the PS.

R-73 As a result of endpoint attachment request, L2SC MUST also be able to receive service attributes, such as QoS parameters, associated with the L2 sessions.

R-74 L2SC MUST be able to establish more than one L2 session associated with the endpoint attachment, e.g. a single authentication can trigger more than one session establishment for example to provide a TLS and IP service for a business customer..

R-75 As the result of authentication, L2SC MUST be able to receive, from PC, a pair of VIDs per L2 session, one on the I-NNI interface (on the Va RP) of a L2A and one on UNI interface of a L2A for each L2 session.

R-76 L2SC MUST be able to setup an IVC by providing the association between the pair of VIDs assigned by PS for each L2 session.

R-77 L2SC MUST be able to monitor the connectivity of L2 sessions and delete the L2 session upon the detection of failure or a L2 session timeout..

R-78 L2SC MUST support the deletion of L2 sessions upon timeout of a session or request from PS. [Ed: for further work on timeout trigger]

The L2SC Function MUST be able to receive/activate/modify/delete QoS policies via PC for L2 sessions in the associated STF Function.

7.6 Requirements for L3SC in the IP Service Layer

R-79 The L3SC Function MUST be able to glean from the data or control plane packets the necessary information to identify the L3 Session

R-80 The L3SC Function MUST be able to control IP Address assignment using local and remote address Pools.

R-81 The L3SC Function MUST be able to interwork with IP Address allocation mechanisms when it is not assigning the IP Address itself.

R-82 The L3SC Function MUST be able to associate QoS policies with L3 sessions and have these instantiated in the associated STF Function

R-83 The L3SC Function MUST be able to dynamically manage the QoS policies in the associated STF in accordance with the L3 Session's characteristics

R-84 The L3SC Function MUST be able to manage Multicast forwarding behaviour in the appropriate associated functional modules in accordance with the multicast architecture

R-85 The L3SC Function MUST be able to associate OAM with L3 session state

R-86 The L3SC Function MUST support AAA client functionality.

R-87 The L3SC Function MUST support the attaching of policies to L3 sessions and MUST support the dynamic changing of policies

R-88 The L3SC Function MUST support the interactions with the PC Function to manage policies

R-89 The L3SC Function MUST support the implementation of the L2 Control Protocol as per TR-147 [5]

R-90 The L3SC Function MUST support profile based configuration and bulk provisioning.

R-91 The L3SC Function MUST support security capabilities to detect and block L3 Denial of Service (DOS) attacks and Theft of Service (spoofing)

R-92 The L3SC Function MUST provide a service redundancy mechanism for service continuity in case of node failure, uplink failure (network facing link), or downlink failure (customer facing link) .

R-93 The L3SC function MUST provide an active/standby election mechanism between L3F functions residing in different nodes.

7.7 Requirements for L3RC in the IP Service Layer

R-94 The L3RC function MUST be able to keep track of the bandwidth of the signalled unicast flows being forwarded across each interface of the L3F function on a per Class of Service basis.

R-95 The L3RC function MUST be able to receive requests from an application through out of band signalling, or network based signalling.

R-96 The L3RC MUST check the available bandwidth specific to the class associated with the application request and respond as to whether or not the request can be honoured.

R-97 The L3RC MUST be able to notify the PC about local interface status changes.

R-98 The L3RC function MUST be able to control how many unique multicast groups the L3F can replicate across its interfaces. The L3RC function MUST be able to do this on a per physical interface, per logical interface and on a group of multicast groups basis

R-99 The L2A function inside the Ethernet Service Layer MUST be able to signal to an upstream L3RC when it starts to replicate each multicast group across a specific UNI interface.

R-100 The L2A function inside the Ethernet Service Layer MUST be able to request from an upstream L3RC whether it is allowed to replicate a certain multicast group across a specific UNI interface.

R-101 The L3RC function MUST be able to correlate the multicast resource requests, the unicast resource requests, and the per user multicast resource requests to allow or disallow resources on the network links.

7.8 Requirements for UNI

R-102 The L2A@UNI at the T reference point MUST support MEF UNI compliant with MEF 10.2.1 [20], MEF 13[21], and MEF20 [22].

R-103 The L2A@UNI MUST support upstream traffic conditioning to accommodate a L1 speed mismatch

7.9 Requirements for E-NNI

R-104 The E-NNI –L2 MUST support MEF 4 [17] , MEF E-NNI Baseline Spec (MEF26.1 [23]).

R-105 The ENNI-L2 SHOULD support the MEF E-NNI Amendment Spec (MEF28 [24]) requirements.

7.10 Requirements for L3F

R-106 The L3F Function MUST support the routing requirements in TR-92.

R-107 The L3F function MUST support IPv4 and IPv6 multicast forwarding.

R-108 The L3F function at the SI-NNI interface inside the IP Service Layer MUST support multicast router discovery using the Multicast Router Discovery Protocol defined in IETF RFC 4286 needed by IGMP/MLD snooping switches.

R-109 The L3F Function MUST provide OAM mechanisms in order to detect node and uplink/downlink failures.

R-110 The L3F Function MUST provide unicast/multicast downstream stream redirection after an active/standby switchover event.

Appendix I - Technology specific functional modules

The following table shows how technology specific modules are derived from technology independent modules before mapping to network nodes.

Technology independent module A		Technology independent module B		Technology independent module C	
Technology specific module Technology A1	Technology specific module Technology A2	Technology specific module Technology B1	Technology specific module Technology B2	Technology specific module Technology C2	Technology specific module Technology C2
Broadband Network Nodes					

Table 2 Mapping technology independent modules to technology specific modules

I.1. Access technology specific Modules

These Modules are instantiated between the Vc and U1 reference points across the access function set. They are defined by other standards organizations such as ITU.T and IEEE and are not repeated here.

I.2. Synchronization Function

Synchronization function **(Synch)** can be instantiated in a variety of technology dependent synchronization modules, according to:

- The type of synchronization to be provided, i.e. frequency/phase/time.
- The methods used for synchronization extraction and for synchronization distribution: synchronous physical layer (only for frequency synch), packet based (frequency / phase / time), GPS (frequency / phase / time). It should be noted that a synchronous physical layer method can also be used in conjunction with a packet based method.
- Synchronization modules based on physical layer methods differ by the physical layer type:, E1, T1, SDH, xDSL, GPON, XG-PON, and Ethernet.

Synchronization modules based on packet based methods differ in some basic characteristics:

- The protocol: IEEE 1588 [27] , NTP [29], differential CES, adaptive CES.
- The clock type: Master, Slave, Boundary, (1588) Transparent.
- Direction of the synchronization packets, either one-way (only for frequency synchronization) or two-way (frequency/phase/time synchronization).

Synchronization modules may be required to interwork between different physical layers or between packet and physical layer methods.

I.3. Aggregation technology specific modules

This section explains:

- How to implement the IVCs using MPLS, Ethernet and other technologies.
- How coordination is achieved between optical protection schemes and packet layer protection schemes
- How IPv4 and IPv6 forwarding is achieved.

I.3.1. Ethernet Adaptation & Forwarding

This functionality would be common to PB and PBB Adaptations. The underlying control plane could be one that is suitable for comparatively trivial networks in the form of G.8032 ring protection, or be solutions suitable to larger Ethernet based aggregation networks, such as 802.1Qay PBB-TE and 802.1aq SPB/SPBB.

The Module associated with the functionality is called L2F-E and performs 802.1 based bridging

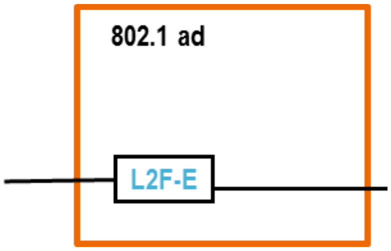

Figure 26 Ethernet Aggregation/Switching

I.3.1.1. Ethernet Provider Bridge Functionality

This functionality would define how an edge device mapped a customer UNI onto an L2VPN instantiation using 802.1ad provider bridging. (This would be directly analogous to the VLAN architectures described in TR-101).

The L2A-E module is performing the needed VLAN encapsulation/removal/addition/translation and the L2F-E Module would be performing the 802.1ad bridging functionality

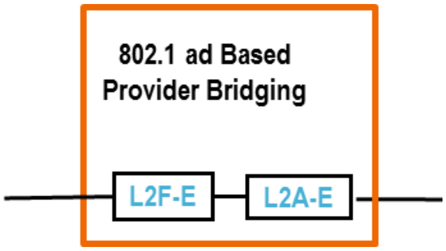

Figure 27 802.1ad based Provider Bridging

I.3.1.2. MPLS L2VPN PE Functionality

In order for the TR-145 architecture to support MPLS L2VPNs, several pieces of functionality would need to be supported on the nodes that implement this architecture. The following list is partial and detailed requirements will be provided for on a node basis in WT-178:

- RFC 3985, 4446,4447, and 4448 for point to point L2VPN
- RFC6073, for Point to Point L2VPN
- RFC 4664, 4761, and 4762, for multipoint L2VPN
- RFC6074, section 3.2.2.1. BGP-based auto-discovery, for LDP signaled VPLS
- The protocols and procedures required to support IP multicast traffic within VPLS, as being defined at the draft-ietf-l2vpn-vpls-mcast.
- The protocols and procedures as described in draft-ietf-l2vpn-pbb-vpls-pe-model- and draft-ietf-l2vpn-vpls-ldp-mac-opt are needed to support the PBB-VPLS functionality.

I.3.1.3. Ethernet Provider Backbone Bridging functionality

This functionality defines how an I-component mapped customer Ethernet traffic onto an Ethernet aggregation module for the MEF 6.1 service set (ELINE, ELAN and E-TREE services) using 802.1ah and 802.1aq Adaptations.
The L2A-E is performing the needed VLAN encapsulation/removal/addition/translation and L2A-B (Provider Backbone bridging) performing the mapping to I-SID and B-VLAN). The L2F-E then performs the Bridging function.

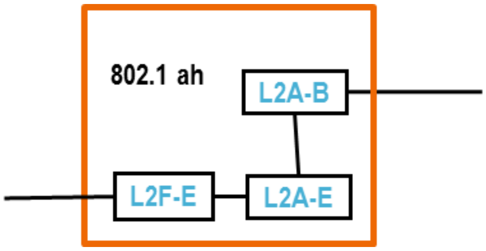

Figure 28 802.1ah based Provider Backbone Bridging

I.3.2. MPLS Adaptation and Forwarding

I.3.2.1. MPLS LSR and LER Functionality

The MPLS specific instantiations map to the L2A and L2F technology independent functional modules.

- L2A-M: generic Adaptation function, providing MPLS Adaptation and encapsulation.
- L2F-M: generic forwarding function, providing MPLS forwarding and processing of forwarding tags.

With some services (e.g. Ethernet transport over MPLS), some Ethernet specific Adaptations and Forwarding functionality can be collocated with the L2A/F-M modules, and they will be referred to as L2A-E and L2F-E.

MPLS Based Aggregation defines the usage of MPLS LSPs signaled by LDP (RFC 5036) or RSVP-TE (RFC3209). This function enables the setup of MPLS LSPs or TE-tunnels between MPLS edge nodes of a network. For LDP signaling, LDP downstream unsolicited label advertisement mode and LDP downstream-on-demand label advertisement mode must be supported as per RFC 5036. LSPs/tunnels build the transport infrastructure that could be used for VPN and IP services.

This LSP could be described as a sequence of routers:

- which begins with an LSR (an "LSP Ingress") that pushes a level m label, all of whose intermediate LSRs make their forwarding decision by label Switching on a level m label,

- which ends (at an "LSP Egress") when a forwarding decision is made by label Switching on a level m-k label, where k>0, or when a forwarding decision is made by "ordinary", non-MPLS forwarding procedures.

A MPLS edge node [RFC 3031] is an MPLS node that connects an MPLS domain with a node which is outside of the domain, either because it does not run MPLS, and/or because it is in a different domain. Note that if an LSR has a neighboring host which is not running MPLS, that LSR is an MPLS edge node.

In the MPLS forwarding paradigm, once a packet is assigned to a FEC (Forwarding Equivalence Class), no further payload classification is done by subsequent routers; all forwarding is driven by the MPLS labels.

Seamless MPLS addresses the possible spanning of connections between different domains but still providing a low complexity implementation on e.g. access nodes. Seamless MPLS LSPs can be achieved via a subset of the L2A-MPLS and L2F-MPLS functionality.

The following figure shows an MPLS LER and LSR Nodal disposition of functional modules (Note: the ‘other Adaptation functions’ could be e.g. Ethernet Adaptation, Circuit Emulation, ATM, Legacy Adaptation, etc). Note that the L2F-M functionality is omitted, as an LER can logically be seen as a standalone entity. In reality every LER will also have LSR functionality.

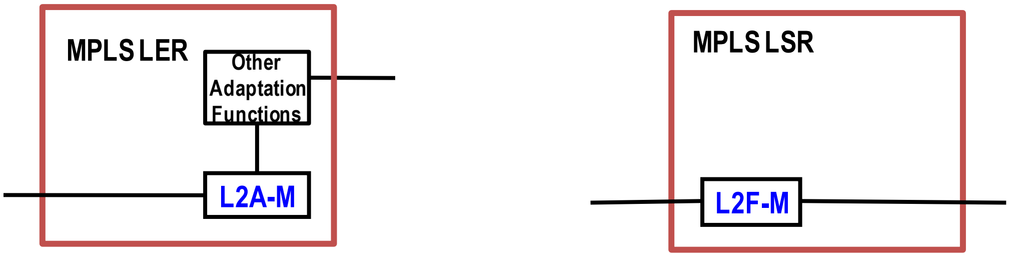

Figure 29 MPLS LSR/LER

Seamless MPLS is a combination of different mechanisms such that certain nodes do not have to support all the IP/MPLS control and data plane requirements while continuing to provide MPLS Pseudowire services. Seamless MPLS also simplifies the setup of LSPs over multiple MPLS domains and provides end to end service independent transport.

From an organizational and operational point of view it may make sense to define the boundaries of such domains along the pre-existing boundaries of aggregation networks and the core network. This means the MPLS LSP starts for example on a node inside the access domain, goes over the core network and will be terminated on a different access domain node.

To ensure a plain control plane on access nodes LDP-extension for inter-area LSPs [RFC 5283] and LDP downstream-on-demand [RFC5036] are used. LDP-extension for inter-area LSPs is required to setup LDP LSPs on a node having only default routing entries to reach its neighbors or aggregated routing entries. LDP downstream-on-demand advertisement mode was originally intended to be used with ATM switch hardware, but there is nothing from a protocol perspective preventing the use in a regular MPLS frame-based environment. In this mode the upstream LSR will explicitly ask the downstream LSR for a label binding for a particular FEC when needed.

In order for the TR-145 architecture to support MPLS, several pieces of functionality would need to be supported on the nodes that implement this architecture. The following list is partial detailed requirements will be provided on a per node basis in WT-178:

- Traffic Engineered Label Switched Paths through RFC4206
- If CAC and sub 100 ms resiliency is required: RSVP signaling through RFC3209
- Point to Point and Point to Multipoint Label Switched Paths, RFC4875 and RFC 6388
- Labeled BGP and RFC3107
- IP routing for the control plane: IGPs such as IS-IS, OSPF, and static routing; IGP-TE (either RFC 3784 ISIS-TE or RFC 3630 OSPF-TE) for TE-LSPs.

I.3.2.1.1. VPWS

The following figure shows how MPLS VPWS can be used to carry traffic across the aggregation network.

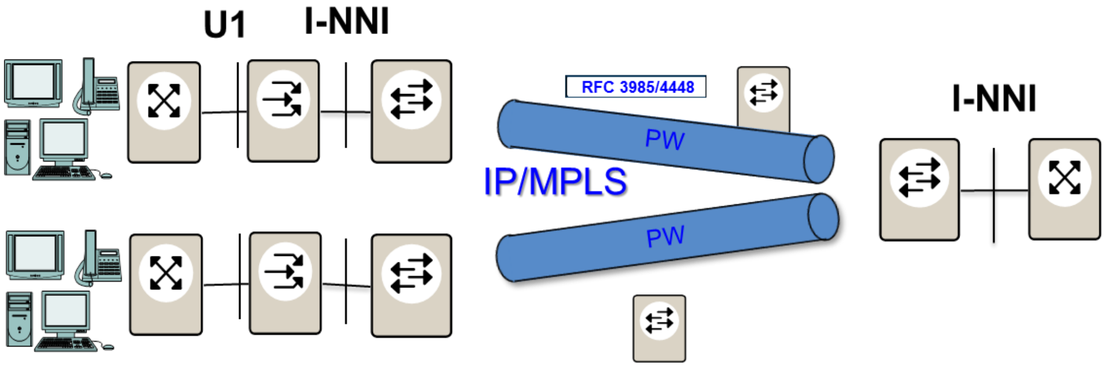

Figure 30 L2VPN – VPWS

When modeling a VPWS capable LER with the functional modules, the Ethernet Adaptation (L2A-E) provides VLAN Adaptation (removal/adding/translating single/untagged/double-tagged traffic), the MPLS adaptation (L2A-M) provide Port and VLAN/EFP based MPLS pseudowire encapsulation, and MPLS Forwarding (L2F-M) is the actual transfer function, as shown in the following picture:

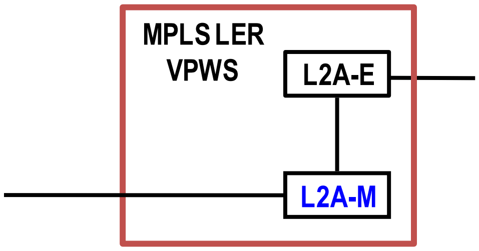

Figure 31 VPWS LER

I.3.2.1.2. VPLS

The following figure shows how MPLS VPLS can be used to carry traffic across the aggregation network.

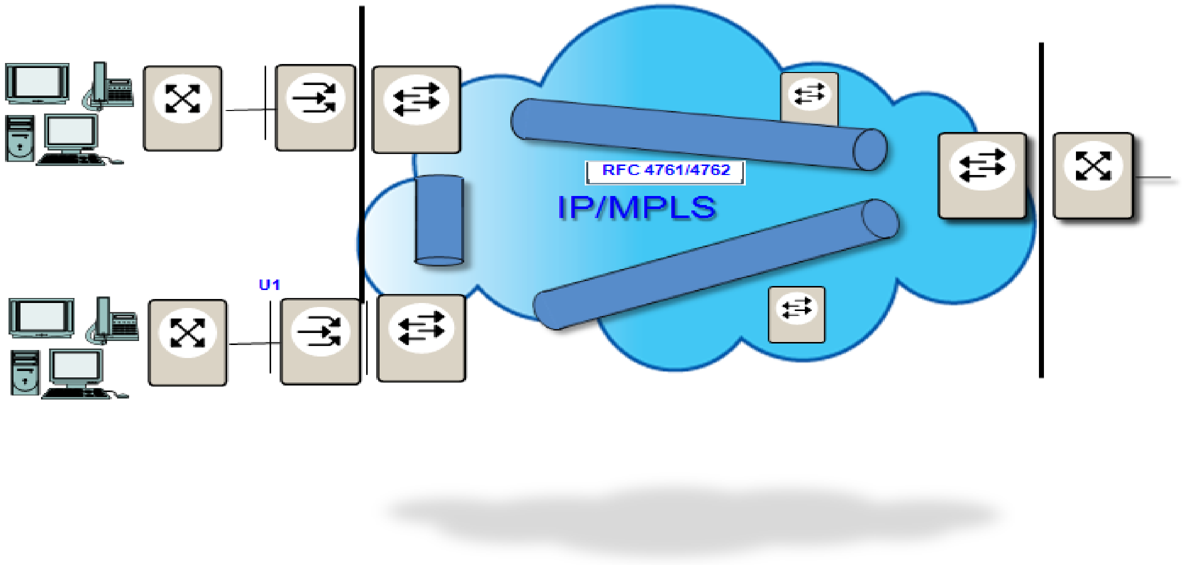

Figure 32 L2VPN – VPLS

When modeling a VPLS capable LER with the functional modules, the Ethernet Adaptation (L2A-E) provides VLAN Adaptation (removal/adding/translating single/untagged/double-tagged traffic) as well as IGMP snooping support, the Ethernet Bridging (L2F-E) provides the VSI functionality of VPLS, the MPLS adaptation (L2A-M) provides mapping from VSI to MPLS Pseudowire and associated encapsulation, and the MPLS Forwarding (L2F-M) provides the NNI connectivity, and as shown in the following picture:

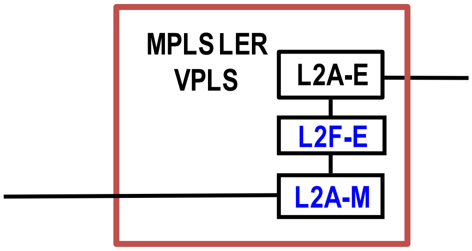

Figure 33 VPLS LER

I.3.2.1.3. MPLS PBB-VPLS Functionality

The following figure shows how PBB-VPLS can be used to carry traffic across the aggregation network

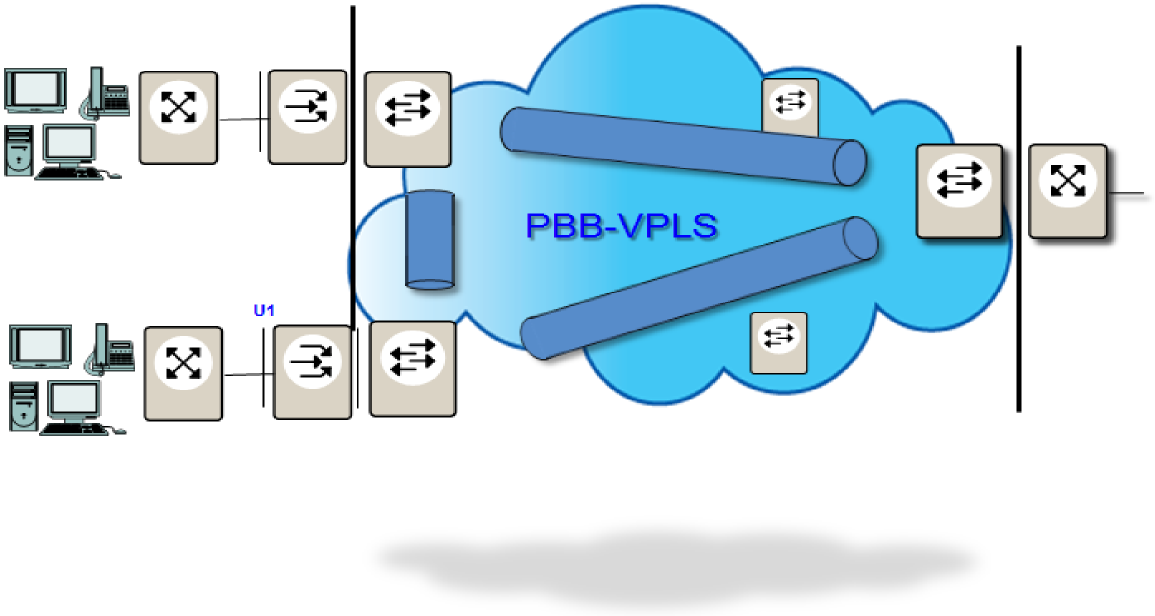

Figure 34 L2VPN – PBB-VPLS

When modeling a PBB-VPLS capable MPLS PE, one needs L2A-E (performing VLAN translation/removal/addition for single/dual/untagged traffic), and additional L2A-B (L2A, Provider Backbone Bridging, adding the I and B-components (I-SID and B-VLAN), an L2F-E (performing Provider Backbone Bridging), and the L2A-M and L2F-M before mentioned to map the B-VLAN to a VPLS instance, as per the below picture.

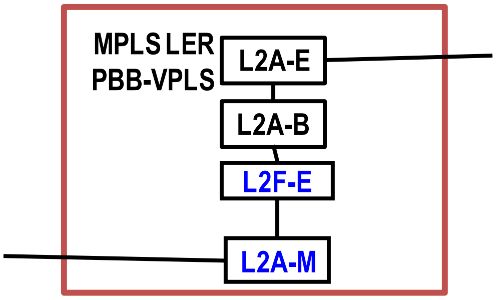

Figure 35 PBB-VPLS functionality

I.3.3. L3 Forwarding

I.3.3.1. IP VPN

I.3.3.1.1. MPLS L3 VPN for IPv4

The following figure shows how BGP/MPLS IP VPNs can be used to carry traffic across the aggregation network.

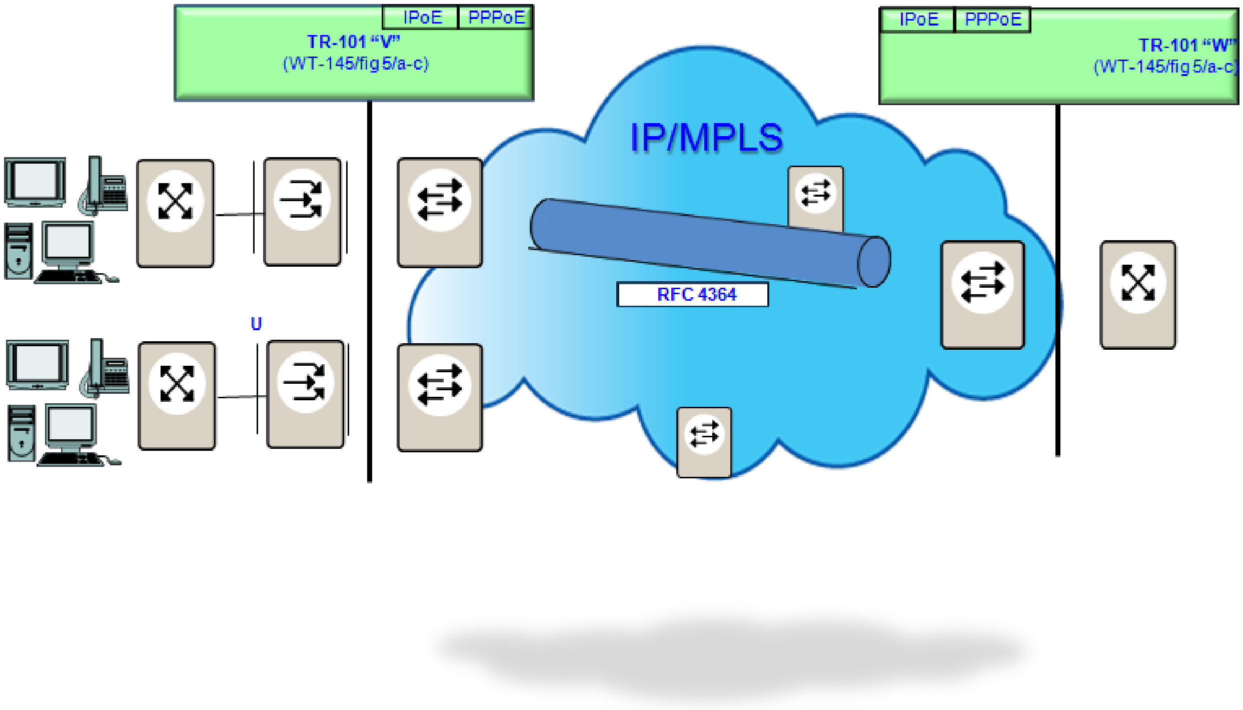

Figure 36 IP/MPLS L3VPN

Modeling an L3VPN capable LER with the functional modules, we need to look at Ethernet Adaptation (performing VLAN termination and ARP functions), IP unicast/multicast Forwarding

(L3F), MPLS adaptation (L2A-M providing mapping from VRF/VR to MPLS encapsulation), MPLS Forwarding (L2F-M), as shown in the following picture:

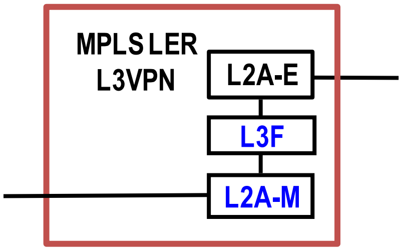

Figure 37 L3VPN LER

In order for the TR-145 architecture to support MPLS L3VPNs, several pieces of functionality would need to be supported on the nodes that implement this architecture. The following list is partial, detailed requirements will be provided for on a node basis in WT-178:

- BGP/MPLS L3VPN as per RFC 4364 and MP-BGP as per RFC2858
- Support the protocols and procedures required to support IP multicast traffic within MPLS L3VPNs as defined in RFC6517 (mandatory minimal set), RFC 6514 (BGP encodings and procedures) and RFC6513 (contains all the options).

I.3.3.1.2. 6vPE MPLS L3VPN Functionality

The following figure shows how BGP/MPLS IPv6 VPNs can be used to carry IPv6 traffic across the aggregation network (MPLS IPv4 or MPLS IPv6 aggregation network).

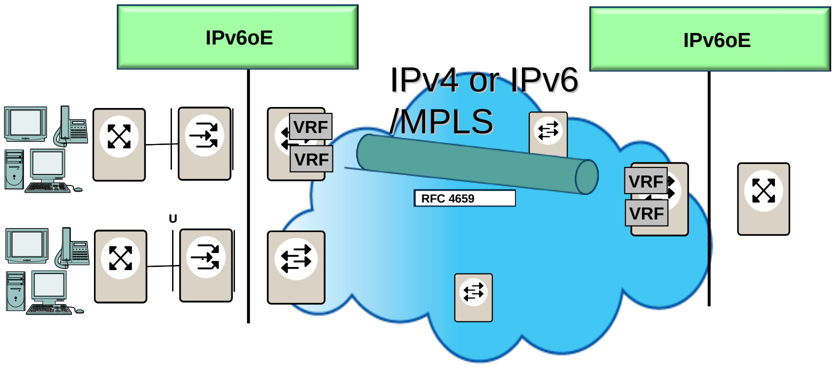

Figure 38 BGP-MPLS based IPv6 VPN

VPN Routing and Forwarding tables (VRFs) maintain the reachability information and forwarding information of each IPv6 VPN separately.

The functional decomposition is the same as before, except now L3F supports the forwarding of IPv6. In case of an MPLS IPv4 aggregation network the L2A-M and L2F-M are still leveraging an IPv4 control plane. In case of an MPLS IPv6 aggregation network the L2A-M and L2F-M are leveraging IPv6.

The functionality required is defined in:

- RFC 4659 – BGP-MPLS IP VPN Extension for IPv6 VPN
- Support the protocols and procedures required to support IP multicast traffic within MPLS L3VPNs as defined in RFC6517 (mandatory minimal set), RFC 6514 (BGP encodings and procedures) and RFC6513 (contains all the options).

I.3.3.2. Flat IP routing

I.3.3.2.1. Flat IP without tunneling

iBGP and at least one IGP (ISIS or OSPF) supporting multiple areas and multiple instances needs to be supported as well as static routes.

The use of iBGP in large IP networks is required to preserve the IGP scalability and performances. iBGP is used to provide a view of the external routes to the intra-AS BGP routers. To guarantee the scalability of BGP, BGP Route Reflectors may be used in the aggregation networks.

In order for the TR-145 architecture to support IP multicast traffic, PIM-ASM, PIM-SSM protocol and SSM translation/mapping needs to be supported.

I.3.3.2.2. 6PE MPLS L3 Functionality

The following figure shows how 6PE approach can be used to carry IPv6 traffic across an MPLS enabled IPv4 aggregation network.

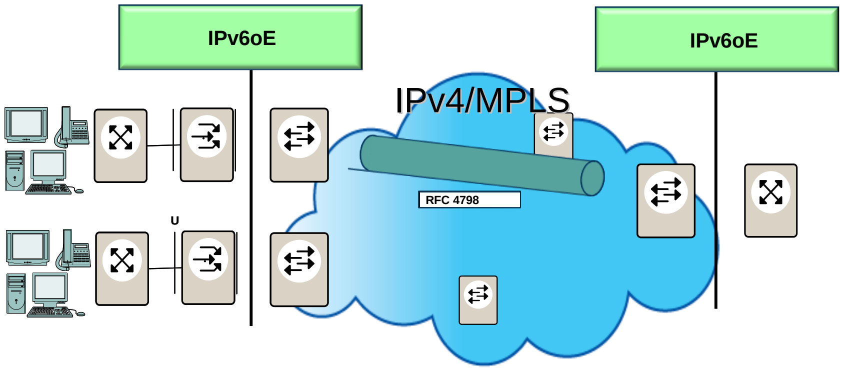

Figure 39 The 6PE approach

The major difference is in the L3F functionality, where now IPv6 is leveraged, although the MPLS constructs (L2A-M and L2F-M) are still leveraging an IPv4 control plane. For further details see RFC 4798 – Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE).

In order for the TR-145 architecture to support IP multicast traffic, several pieces of functionality would need to be supported on the nodes that implement this architecture:support the protocols and procedures required as being defined in draft-ietf-mpls-mldp-in-band-signaling.

Appendix II - Use Cases

II.1. Leveraging underlying OAM schemes

Adaptation to a lower layer PHY layer (e.g. G.709/WDM) can happen within a discrete, dedicated node of the network, through the L1A function. The L1A function performs adaptation to the lower layer (e.g. using a digital wrapper like G.709), together with OAM and other supporting functionality between itself and its peer L1A at the other side of the (virtual) link.

In case the OAM functionality notices a continuity problem, it needs to convert this event to the appropriate signals towards the higher layer functions such they can react appropriately. Figure 40 shows a certain disposition where IP packets are adapted onto Ethernet (through the L2A$_a$ and L2A$_d$ functions.) Other L2A's at the other side of the wire hand-off Ethernet packets to lower layer L1A functions (and example could be WDM G.709 digital wrapper functionality).

If the L1A associated OAM functions detect a loss of continuity, the L2A's colocated with the L1As (L2A$_b$ and L2A$_c$ respectively) have to signal the peer L2As (L2A$_a$ and L2A$_d$ respectively) through sending e.g. a Loss of Signal.

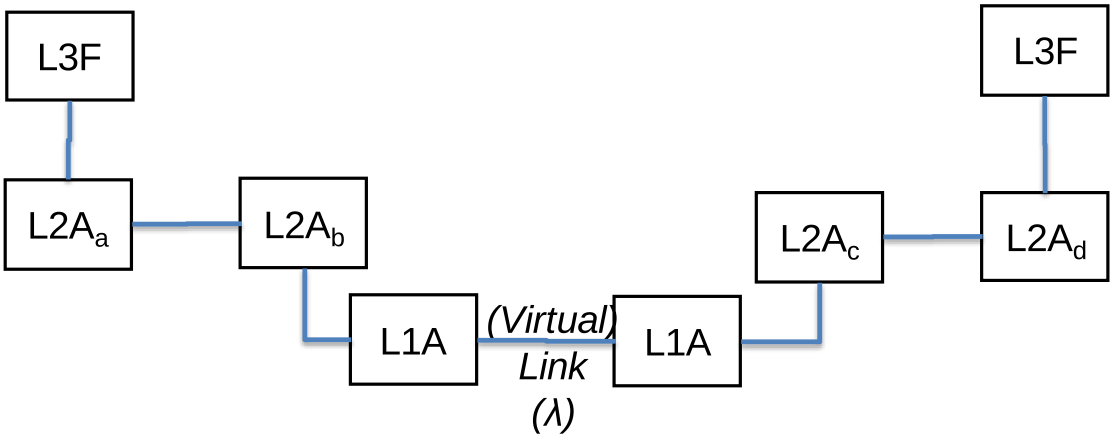

Figure 40 OAM functionality conversion across diverse L2A's

Another approach could be to have all functions collocated (L3F, L2A and L1A), as shown in Figure 41. The advantage of this approach is that the L2A function can be made aware of the OAM signals and other functionality happening at the lower layers much sooner.

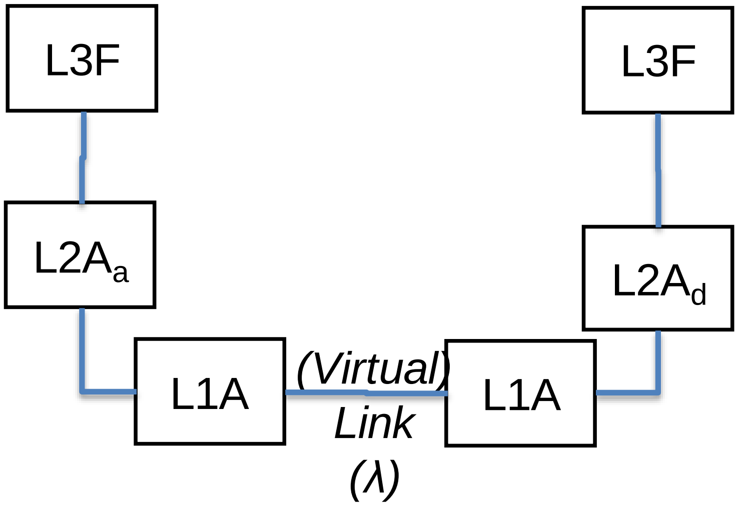

Figure 41 OAM functionality with L2A and L1A collocated

II.2. Multiservice support through pseudowires

TR-145 supports multi-services on common network infrastructure. This section describes how the concept of pseudowires (PW) can be leveraged and extended to provide as a logical connectivity instance to support multi-services on the common platform.

- PW can be used in access, aggregation, and the core network to provide multi-service transport.
- A pseudowire does not necessarily require the use of an underlying MPLS transport layer, in particular when pseudowires are extended across an Ethernet access network to the CPE.
- On the Aggregation and core network use VPLS/VPWS PW on per service basis
- Mapping between Access S-VLANs and tunnel LSPs on the Aggregation and core network while performing pseudowire stitching.

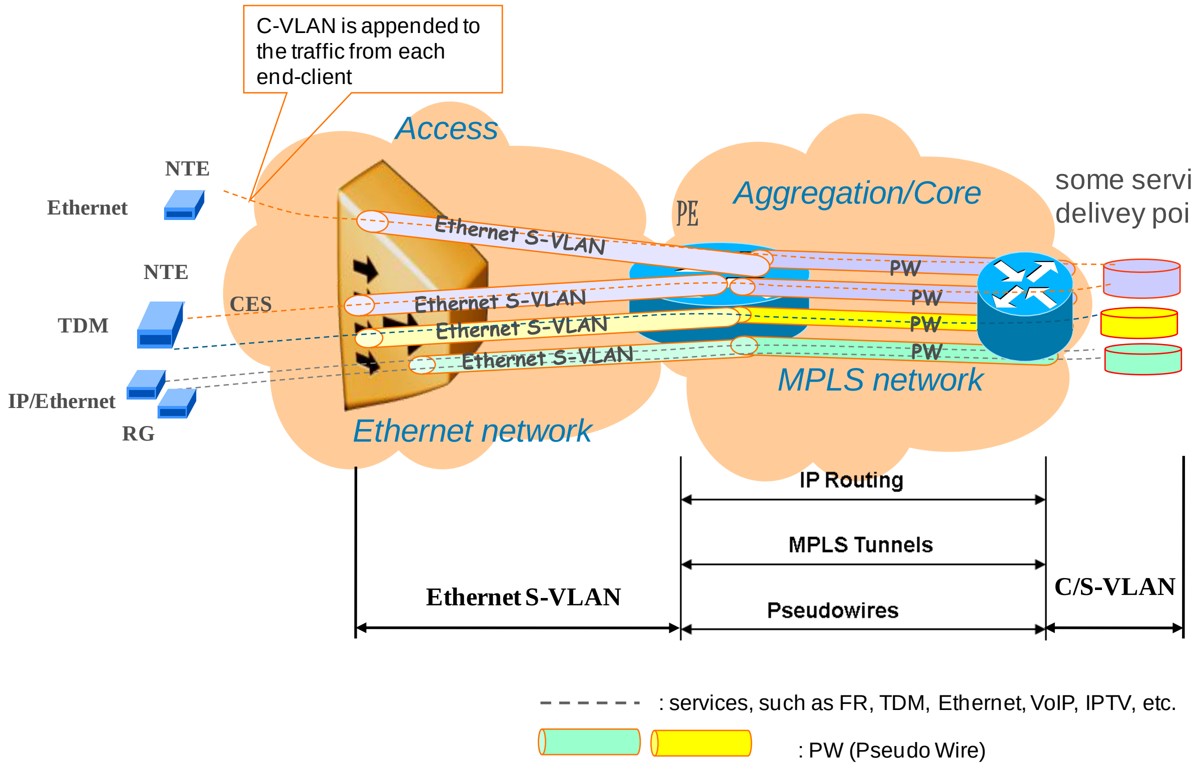

Figure 42 Ethernet maps to MPLS PW

Figure 42 shows an example of using extended PW, i.e. using Ethernet VLAN structure and PW in the aggregation and core network to deliver multiple types of service across the common infrastructure. Here the S-VLAN is used identify each service in the Ethernet network. The S-VLAN is further mapped to a PW in the MPLS network, where the S-Tag can be removed or encapsulated within the PW payload.

II.3. Establishing connectivity for Multi-service endpoints to a multi-service network with a common access

This section provides the use cases for establishing connectivity for multi-service endpoints to a multi-service network with a unified access, using the session control function. The endpoint devices can be legacy devices (e.g. POTS), Ethernet service devices, IP service devices (both v4 and v6) or Application service devices (e.g. hosts).

Figure 43 depicts how L2SC be used to establish Ethernet connectivity from an endpoint at UNI-L2 to an service edge with an E-NNI-L2 interface. When the Ethernet endpoint is brought into service and requests attachment to the access network in order to connect to one or more service edges, the L2SC performs access session control (e.g. 802.1X, red arrows) and instructs the Layer 2 Adaptation (L2A) functions to establish the logical connectivity (green arrow) between the UNI-L2 and the existing IVC (black arrow in the drawing) and thus towards the E-NNI-L2.

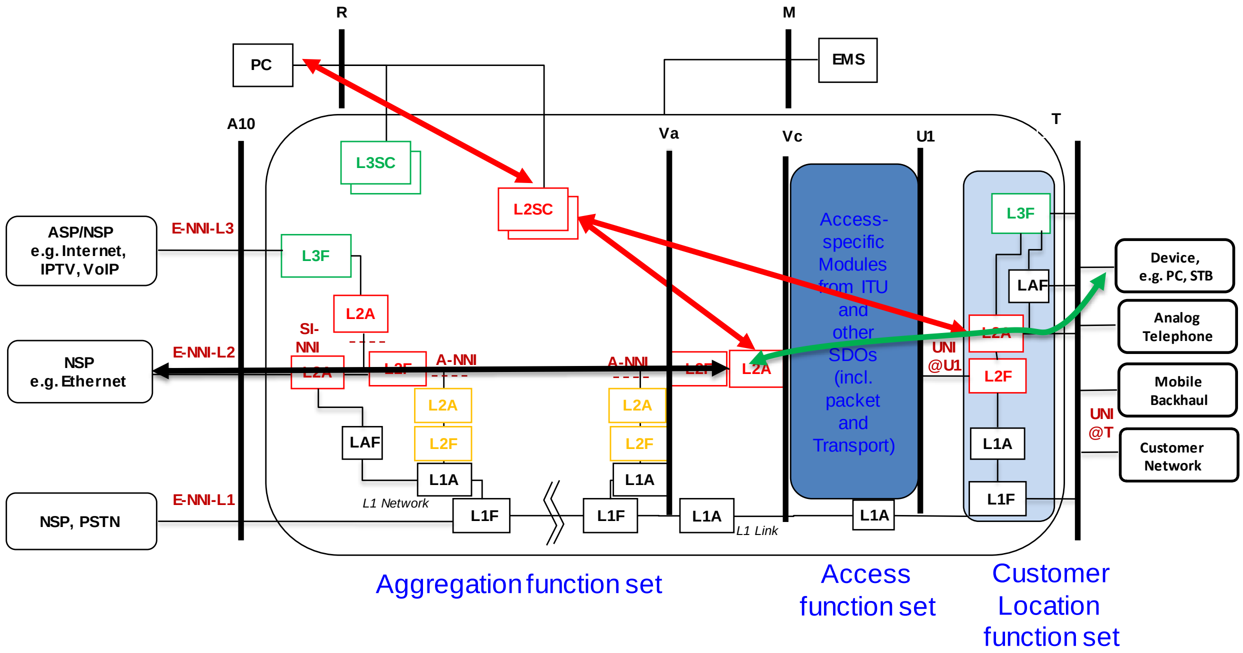

Figure 43 Establishing connectivity for Ethernet endpoint to an Ethernet service network

Figure 44 depicts how L2SC/L3SC be used to establish connectivity from a legacy endpoint, such as POTS or a PBX to an IP service infrastructure leveraging VoIP and then be further delivered to a PSTN network or an NG packet based service network (yellow arrow). The session establishment happens in two steps. The first is similar to the case where the endpoint is an Ethernet as presented in Figure 43, except that the LAF initiates the session establishment on behalf of a legacy device (green arrow) and L2SC establishes the logical connectivity to the LAF at the customer location by configuring the appropriate L2A functions (red arrows). A 2nd step (e.g. when using VoIP to transport POTS) provides the service level authentication and session connectivity between LAF and service provider network using L3SC (orange arrows).

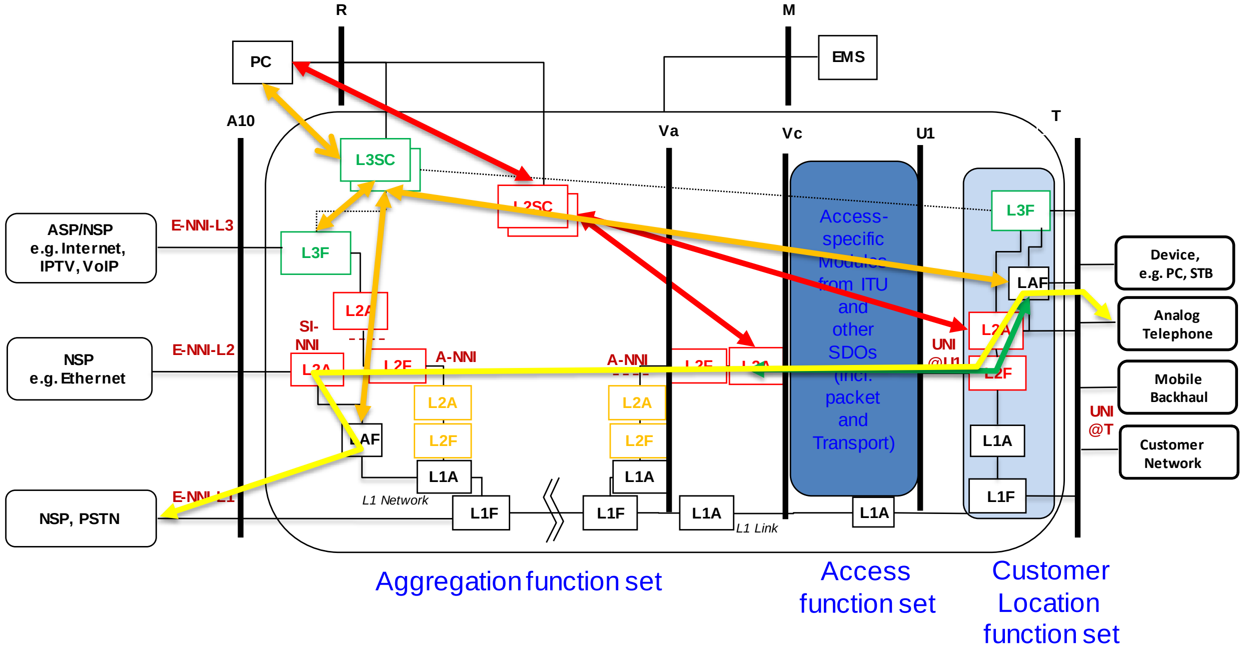

Figure 44 Establishing connectivity from a legacy service endpoint to an IP service network

Figure 45 depicts how the session control function be used to establish connectivity from an IP endpoint to an IP service provider network. The session establishment happens again in two steps, similar to Figure 44, it relies on some form of authentication (red arrow) to establish Ethernet connectivity (green arrow), and in step 2 establish an IP session with the IP service edge(s) (yellow arrow).

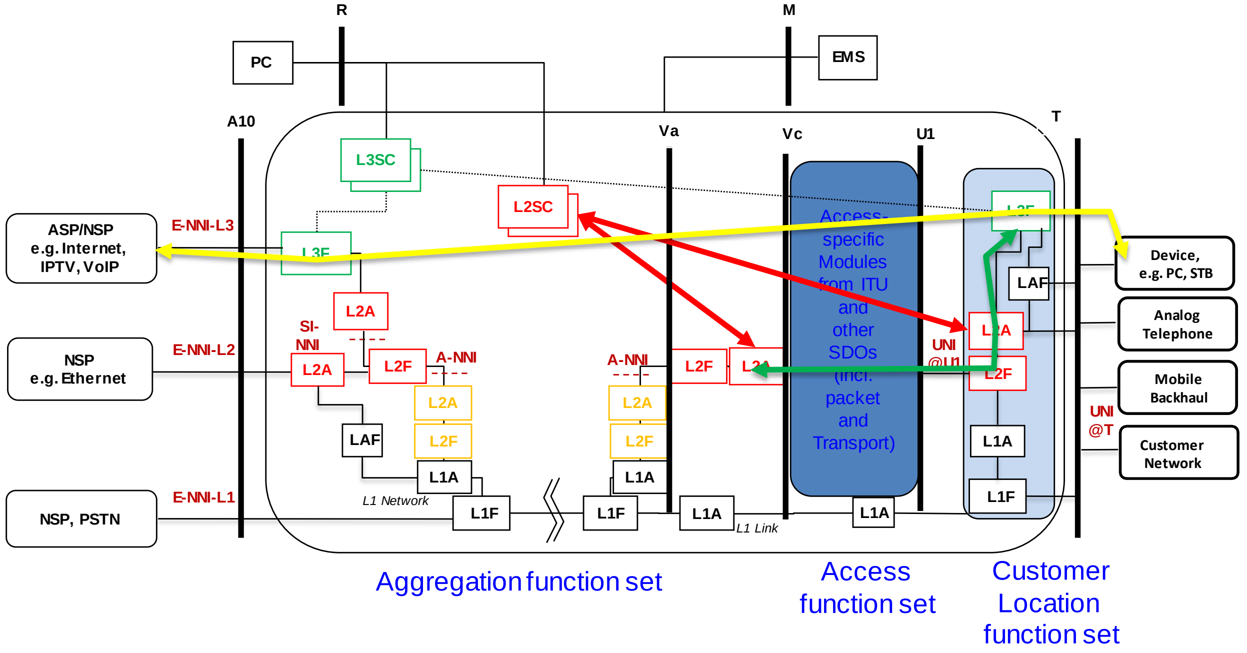

Figure 45 Establishing connectivity from an IP endpoint to an IP service network

Note that more than one access session can be established with a single L2SC interaction. For example, a business router can use a single authentication to establish access to both a TLS and ISP from the same access. If the same SP provides both Ethernet access as well as the TLS or IP services, then authentication credentials need not be exchanged twice. However, in the general case, there need to be separate mechanisms, one for allowing the establishment of establishing L2 connectivity, and one for allowing service association at the service edge.

II.4. Access Service Assurance

Use cases of devices accessing various network services through a common access network have been described in Section II.3. Once connectivity has been established, service assurance and business models may require determination of a failed or reset connection. There is also a need to provide per session data (e.g. the amount of data carried). The required functionality in TR-145 is that once a device is attached to the network and connected to one or more service edges it must be possible to assure that the session remains established and to collect session statistics, including uptime, errors, and the volume of data carried. At a functional level, this means that there needs to be an Ethernet link assurance mechanism, like the OAM developed in TR-101 and TR-156. There also needs to be a measurement and data collection mechanism to determine utilization, and since TR-145 is a multi-service-edge architecture, this will need to be UNI interface.

This applies to the use-cases described in the previous section II.3, as shown in Figure 43, Figure 44 and Figure 45. Service Assurance can for example be deployed to verify continuity on the access specific sessions (the green arrows) or the overall sessions (the yellow arrows).

Note that more than one access can be assured with a single session. For example, the access network can use a single instrumentation event to set up or tear down the access sessions for both a TLS and ISP for the same access. If the same Service Provider provides both Ethernet access as well as the TLS or IP services, then session statistics need not be collected twice. However, in the general case, there need to be two mechanisms that allow assuring access connectivity, and then service association at the service edge.

II.5. Service Redundancy

II.5.1. Overview and motivation

In broadband networks, link/node protection is not enough to provide the required session/service redundancy of all types of subscribers. The subscriber needs to re-establish a new session if a link/node fault happens in the network. Sometimes it is hard for the terminals to automatically perform the session re-establishment procedure, e.g. the terminals which use DHCP to get the IP address (see more information in WT-146). Even though some of the terminals could handle this action, e.g. PPP subscribers, the failover time is far from the requirements of real-time services and high value-added services, such as video service, VPN service, VoIP service, etc.

Since the BNG usually handles a large number of broadband subscribers, it is one of the most important nodes in the broadband network. This use case explains how to provide session/service layer redundancy for the broadband subscribers when the BNG suffers a node or uplink/downlink failure

II.5.2. Service Redundancy Functionality

The service redundancy mechanism includes several parts:

Election Function: This function is responsible for active/standby BNG election. It could be implemented in the L3SC functional module (e.g. VRRP/RFC5798, etc.), see Section 5.1.2.3.

Detection Function: This function is responsible for remote link or node failure detection, including BNG node failure and uplink/downlink link failure. All of these could be implemented in the L3F functional module, see Section 5.1.1.5 (e.g. BFD/RFC5880) or the L2F functional module, see Section 5.1.1.4 (e.g. Ethernet OAM/802.1ag/Y1713) inside Service Edges, Aggregation and Access nodes as appropriate.

Switchover Function: This function is triggered by fault detection and election mechanism (automatic or manual). It is used for Active/Standby Service Edge switchover and includes upstream and downstream flow redirection. The upstream flow redirection function could be implemented in the L2A functional module, see Section 5.1.1.3 (e.g. gratuitous ARP). The downstream flow redirection could be implemented in the L3F functional module for unicast (route refreshing for subscribers after switchover) and also multicast, see Section 5.1.1.5 (re-instantiating multicast state after switchover).

Note that the manual switchover could be used to software upgrade the current active BNG.

State synchronization Function: an optional function which is used for synchronizing session service info (dynamic information such as IP address, MAC address, PPPoE session ID, DHCP lease time, VPN info, multicast info, QoS info etc.) and also configuration information.
Before a node or remote ink failure happens, the Service Edges use the election function to elect an active node, which is responsible for session establishing and subscriber flow forwarding. When a node and or remote link failure happens, the detection function will become aware of that failure, and then trigger the switchover function to do an upstream/downstream redirect. If the synchronization function is implemented for the subscriber session, the subscribers do not need to re-establish the session, since they will not be aware of the failure (the detailed sync method is out of the scope of this document). After a remote link or node failure recovery there is a choice between either automatically switching back, or manually switching back at a more appropriate time.

II.5.3. Service Edge Redundancy Modes

There are two kinds of redundancy protection, 1:1 and 1:N:

- In the 1:1 mode a standby service edge is available for a single active service edge
- In the 1:N mode a standby service edge is available to multiple active service Edges.

II.6. Wholesale and Business Services through L2VPN

It is possible to provide L2 NSP Wholesale with or without dynamic configuration of Ethernet sessions. Use Case II.6.1 describes the static configuration, while use case II.6.2 describes the dynamic configuration.

II.6.1. Wholesale through L2VPN, static configuration

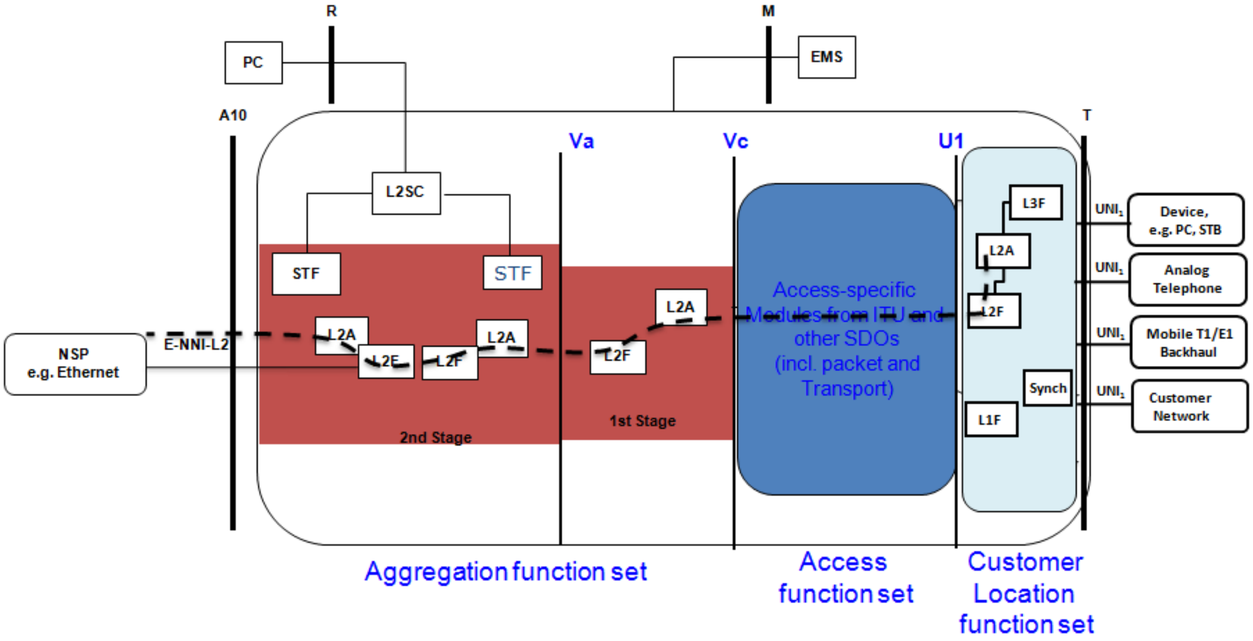

Figure 46 Wholesale through L2VPNs

- Native Ethernet in first aggregation set
- L2A in 2nd aggregation performs L2VPN functionality with SC to auto-provision VLAN to IVC mapping
- IVC between I-NNI and E-NNI (……)
- Native Ethernet handoff
- Alternatively the IVC can be null, allowing for a handoff of the traffic at Va

II.6.2. Wholesale through dynamic configuration of Ethernet sessions

This use cases describes how the L2 connectivity between the UNI and ENNI-L2 can be set-up in a dynamic fashion, triggered by data from a customer. This functionality can be used to automate wholesale access.

II.6.2.1. Session Set-Up

The following describes the events to create state at L2 NSP session set-up.

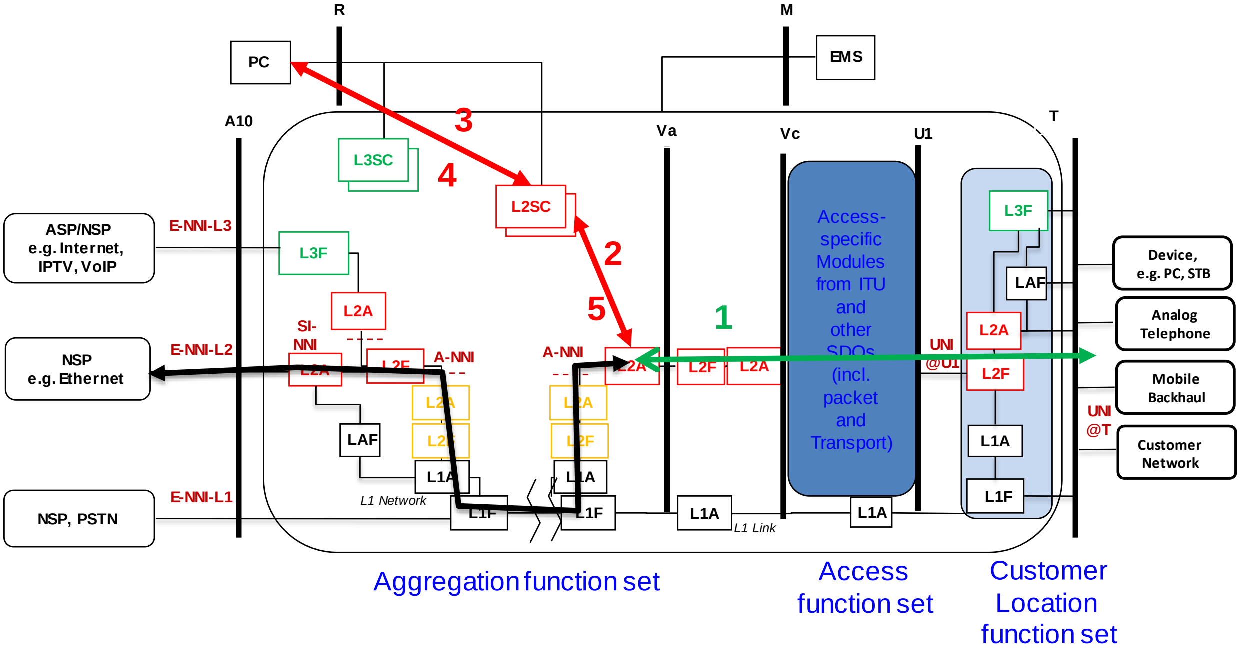

Figure 47 Session Set-up

1. First packet (green arrow).
2. Unique, new (SVLAN, CVLAN) combination is acted upon as a First Sign of Life (FSOL). Any upper layer protocol (IPv4, IPv6, PPPoE, IPoE, MPLS etc.) will be ignored such that there are no dependencies or differing mode of operation.
3. L2SC issues Access-Request to PC/AAA (red arrow)
   Alternatively local or remote policy interface can be used.
4. PC/AAA responds with Accept or Reject, with service attributes. (red arrow)
   Or alternative local or remote policy interaction.
5. Data plane (e.g. forwarding, HQoS etc.) rules instantiated towards a pre-established IVC (black arrow) or can create a recursive L2SC interaction, see use case 0

II.6.2.2. Dynamic Cross-Connect and Backhaul

Once a session is established as in the previous description, traffic is cross-connected to a backhaul path dedicated to the NSP. The backhaul may be across the Regional Aggregation Network, utilizing an IVC. Alternatively, the traffic may be cross-connected to a local E-NNI.

There is an independent backhaul service per NSP, Also the IVC will have the ability to transport many (if not all) access sessions from a particular aggregation node.
The IVC needs to have the following specific attributes.

- supporting many (S-VLAN, C-VLAN) sessions
- Point-to-Point connectivity to the E-NNI interface i.e., no forwarding based on MAC SA/DA
- Per IVC OAM

There can be a recursive L2SC interaction as a result of the steps outlined in the previous use case, where a policy is pulled (steps 6 and 7, red arrows, Figure 48) to active, provision the backhaul IVC (step 8, yellow arrow, Figure 48), and stitch it to the EFP that was created by the FSOL in the previous use-case. .

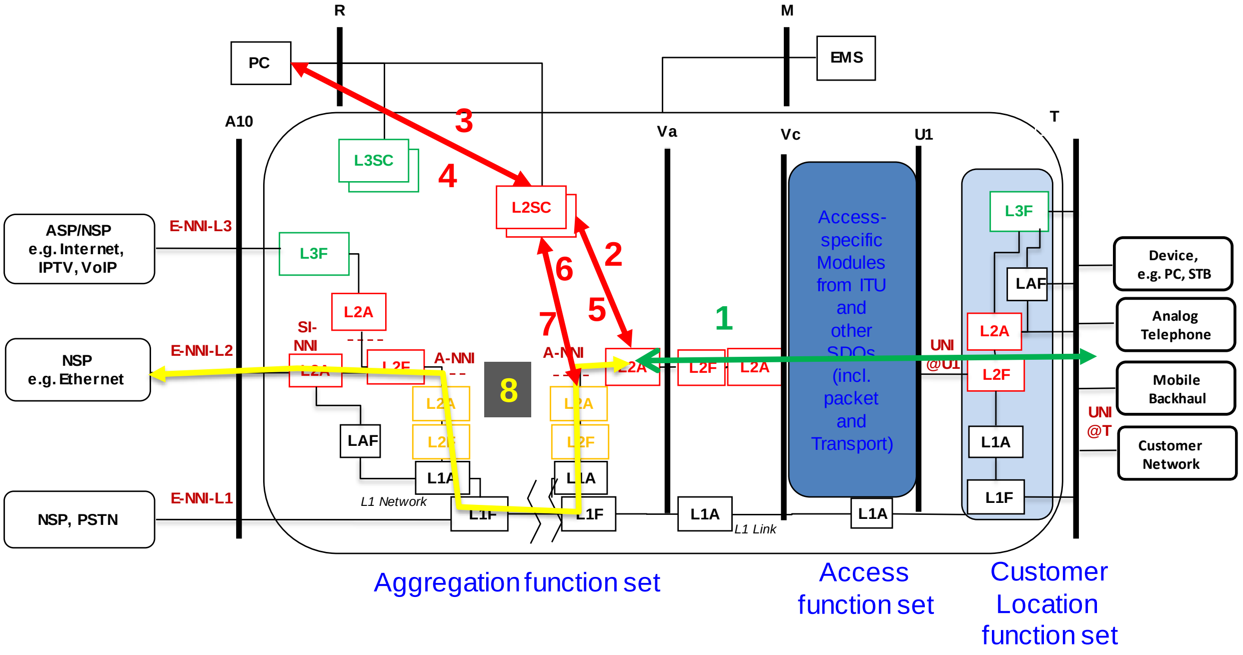

Figure 48 Dynamic Cross-Connect

II.6.2.3. Session Tear Down

The session in the previous use cases can be timed out by either an inactivity timer (as per TR-101 §8.2 R-382) or the session can be terminated from the AAA or Policy Layer, causing the data plane policies to return to their default state.

II.6.2.4. L2 NSP Wholesale Multicast

The L2 NSP Wholesale multicast solution is based on the solution described in TR-147.

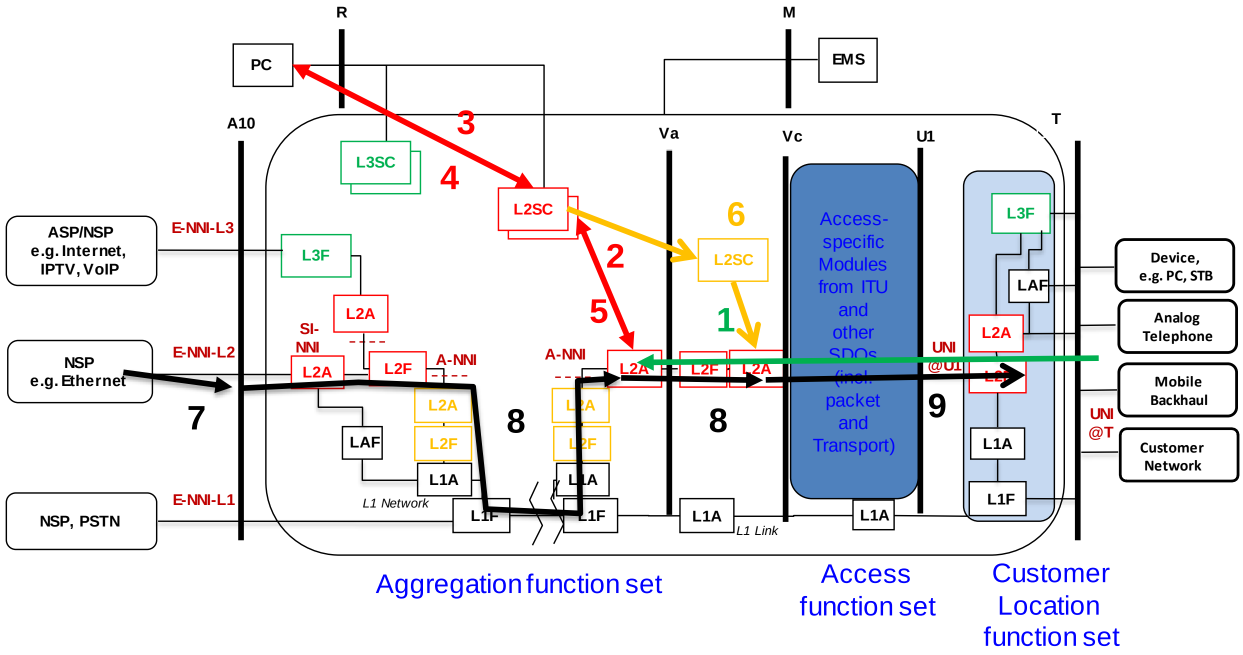

Figure 49 Dynamic Ethernet Sessions with Multicast

Steps 1-5 are identical to use case II.6.2.1, except that in step 4 with an additional service attributes “Multicast Domain”. Alternatively per-subscriber session attributes can be stored locally to the L2SC. The following steps occur:

6. ANCP signals association of a given port to a pre-configured MC-VLAN
7. Multicast traffic injected into MSBN in accordance with business rules and policed by associated input policy at E-NNI interface
8. Multicast traffic replicated as per Regional Aggregation Network policy
9. Multicast traffic replicated in accordance with “Multicast Domain” to access line mapping populated in (6) & IGMP snooping state

II.6.3. Automated Provisioning of IVCs as a result of provisioning of other IVCs

Use case II.6.2 presented a model where the Supporting Aggregation Layer automatically stitches a 'new' IVC originating at a UNI in the Ethernet Service Layer, to an IVC across the Supporting Aggregation Layer, based on first-sign-of-life packets (VLAN id's etc). The following use-case does the reverse: if an IVC is provisioned inside the Supporting Aggregation layer, part of that configuration also holds state about the (not yet) provisioned IVCs starting in the Ethernet Service Layer. In other words, this use case allows the network to automatically configure L2 forwarding state (via a provisioning interface) in a certain part of the network, as a result of configuration in a more aggregated part of the network

The context of the use case is as follows and is shown in Figure 50:

- Assumes the Ethernet Service Layer (red) adapted via the Supporting Aggregation Layer (orange)
- Configuring an IVC between the I-NNI@va and the SI-NNI, or the I-NNI@Va and the ENNI
- Configure an IVC between the UNI and the I-NNI@va
- Stitch those two IVC's together
- This use case will automate a large step of these three individual parameters

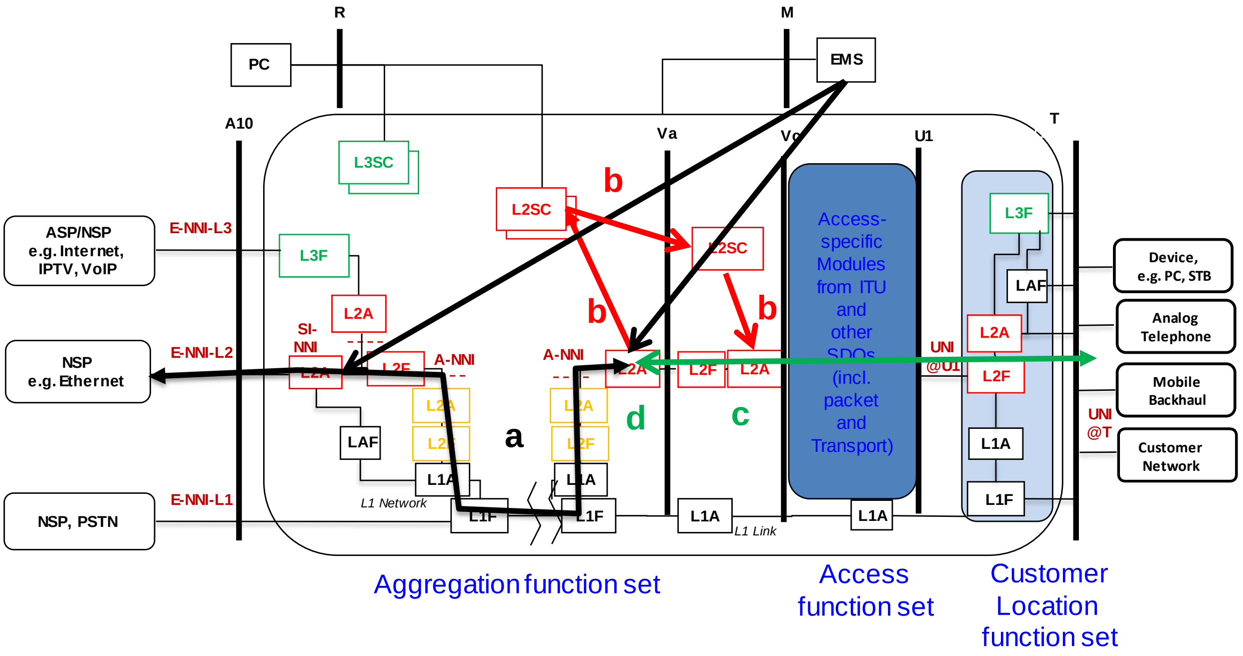

Figure 50 Auto-provisioning IVCs through L2SC to LSC communication

Steps:

a) The IVC is created across the Supporting Aggregation layer (via EMS or through configuring) between the L2A at the I-NNI@Va and the L2A at E-NNI-L2@A10 or SI-NNI interface. (black arrow)

b) This configuration triggers the associated L2SC to signal a provisioning command to the L2A(s) in Ethernet Service Layer (red arrows) (Note : extra information like location and reachability information for this L2A might need to be configured as well as part of the IVC configuration
c) The L2A in Ethernet Service Layer is provisioned to setup another IVC between I-NNI and UNI (green arrow)
d) The two IVCs are stitched together

A real life example of this use-case: an EoMPLS tunnels is set up between the E-NNI-L2 and the I-NNI@Va (between two MPLS PE in other words). As a result of this the PE facing the UNIs will use ANCP to signal the downstream access node to assign a specific port to a VLAN. This VLAN is cross-connected to the EoMPLS tunnel, creating end to end connectivity between the UNI on the DSLAM and the E-NNI-L2 at the remote PE.

II.6.4. Carrier Ethernet Service delivered through Access MEN

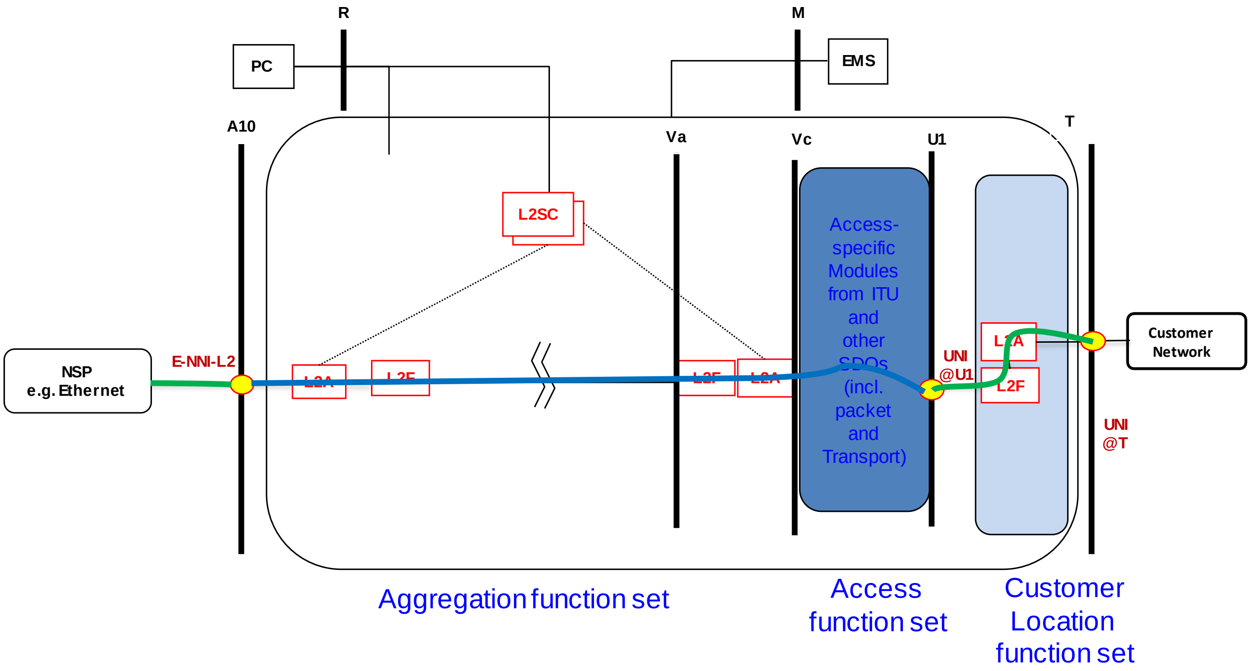

Figure 51 Carrier Ethernet Service Delivered Through MEN

In this use case, the Ethernet service is delivered to a subscriber via an MEF-compliant UNI and MEF-compliant E-NNI. The access network provides a simple Point to Point connection segment of an IVC. Note that the Ethernet service provider supplies and manages the Network Interface Device (NID), which performs the VLAN mapping, 802.1ag/Y.1731 OAM for FM and PM (including SLA instrumentation), and traffic management, etc.

II.6.5. Carrier Ethernet Service delivered through multiple interconnected MEN

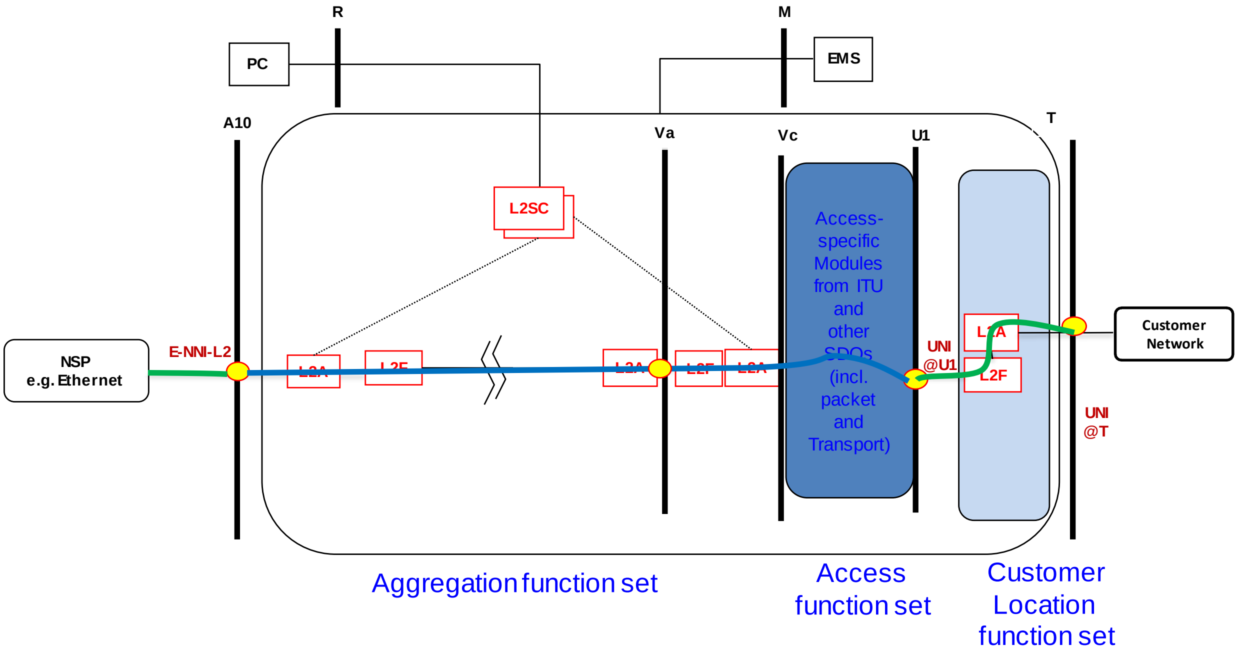

Figure 52 Ethernet Service Delivered through Regional MEN and Access MEN

In this use case, the Carrier Ethernet service is delivered to a subscriber via an access MEN as well as a Regional MEN, i.e. between MEF UNI and the MEF ENNI at the E-NNI-L2 interface. The access network provides a simple Pt-Pt connection segment of an IVC through access network between MEF UNI and MEF E-NNI, at reference point Va, in this case. Any type of network service (e.g. mpt) is delivered through Regional MEN between E-NNI at Va and E-NNI at E-NNI-L2.

II.7. Mobile Backhaul Services

This set of use cases focuses on mobile backhaul, including 2G, 3G, and 4G mobile backhaul. These use cases are also supported in TR-221.

II.7.1. Mobile Backhaul with TDM or ATM service

This section describes the use case of 2G/3G mobile backhaul in which the MSBN provides a T1/E1 or ATM service to the base station/nodeB and mobile operator via Legacy Adaptation Function (LAF), as shown in Figure 53. LAF is used to encapsulate T1/E1 or ATM traffic into Pseudo Wire, which carries T1/E1 or ATM traffic between the pair of LAFs. It provides T1/E1 or ATM hand-off to the Mobile Operator across the E-NNI.

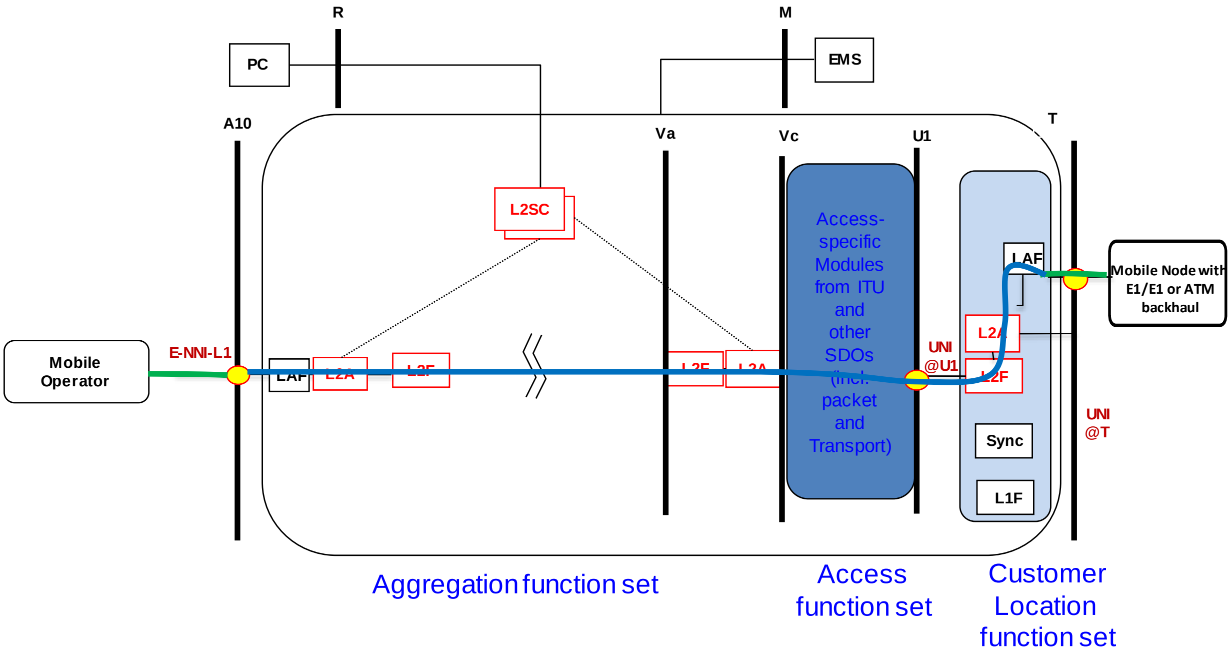
R
M
PC
EMS
A10
Va
Vc
U1
T
L2SC
Access-specific Modules from ITU and other SDOs (incl. packet and Transport)
LAF
Mobile Node with E1/E1 or ATM backhaul
Mobile Operator
E-NNI-L1
LAF
L2A
L2F
L2F
L2A
UNI @U1
L2A
L2F
UNI @T
Sync
L1F
Aggregation function set
Access function set
Customer Location function set

Figure 53 2G/3G Mobile Backhaul – with TDM or ATM interface to Mobile Operator and Cellsites

II.7.2. Mobile Backhaul with Ethernet service

This section describes the use case of 3G/4G mobile backhaul in which an Ethernet service is provided to the nodeB/eNodeB and mobile operator.

Figure 54 presents a use case of Ethernet based mobile backhaul, where the access network provides a simple Pt-Pt IVC and handoff Ethernet to the Mobile Operator across the E-NNI.

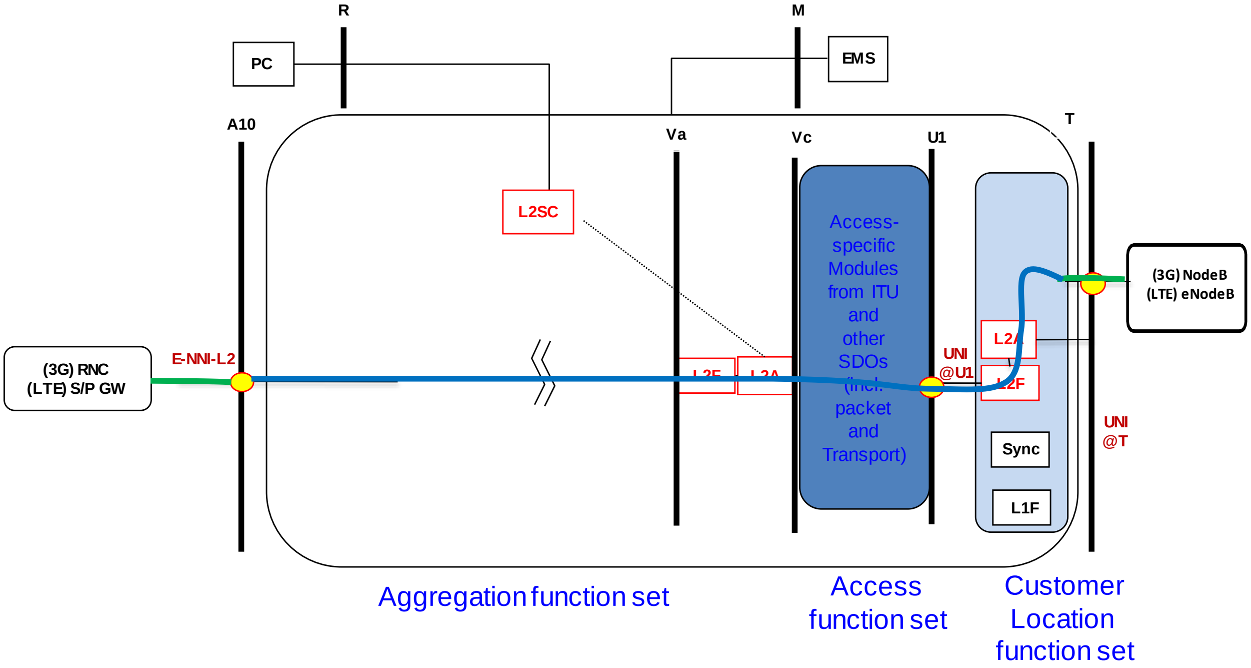

Figure 54 3G/4G Ethernet based Mobile Backhaul with Single Segment IVC via Access

Figure 55 presents a use case of Ethernet based mobile backhaul, where the access network provides a simple Pt-Pt connection segment of an IVC between MEF UNI and MEF E-NNI at Va and E-Line/E-Tree/E-LAN services from E-NNI at Va through MEN network and hands off Ethernet to the Mobile Operator across E-NNI at A10.

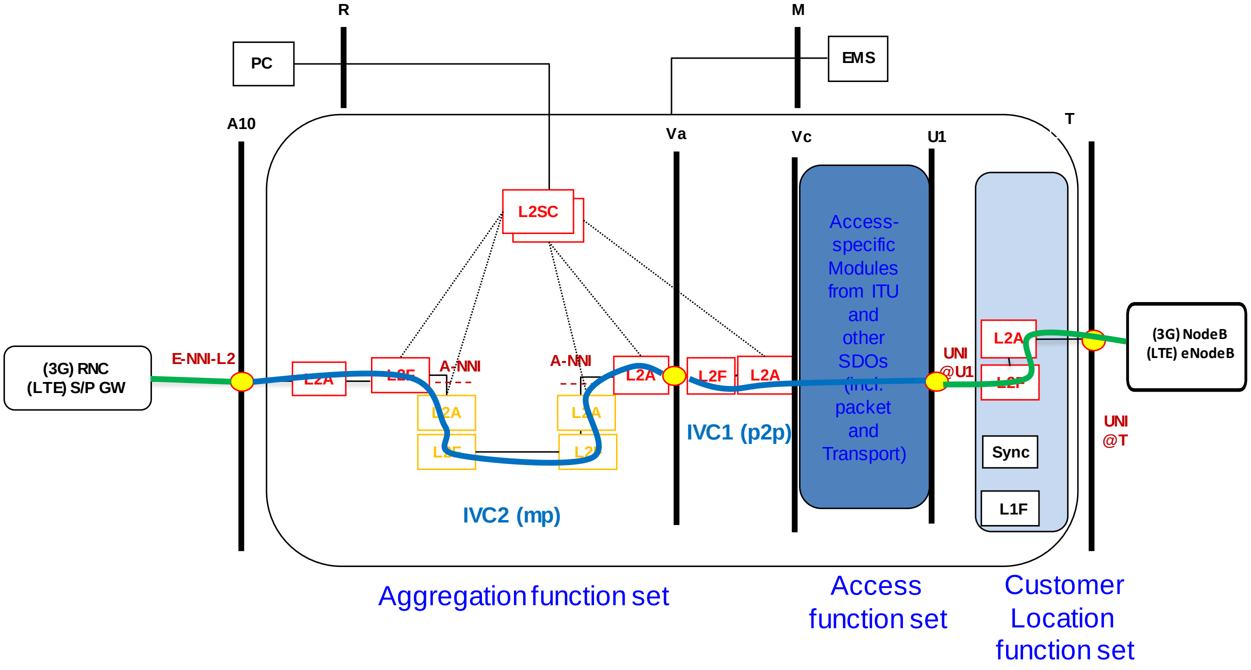

Figure 55 3G/4G Ethernet based Mobile Backhaul with Concatenated IVCs Segments

II.8. Residential Services Use Cases

This set of use cases focus on residential services, such as internet access, IPTV, and Voice. CPE and Home equipment are not shown. L1 services and interworking functional modules are not depicted since they are not relevant for these use cases.

II.8.1. Residential Internet Access

Option 1

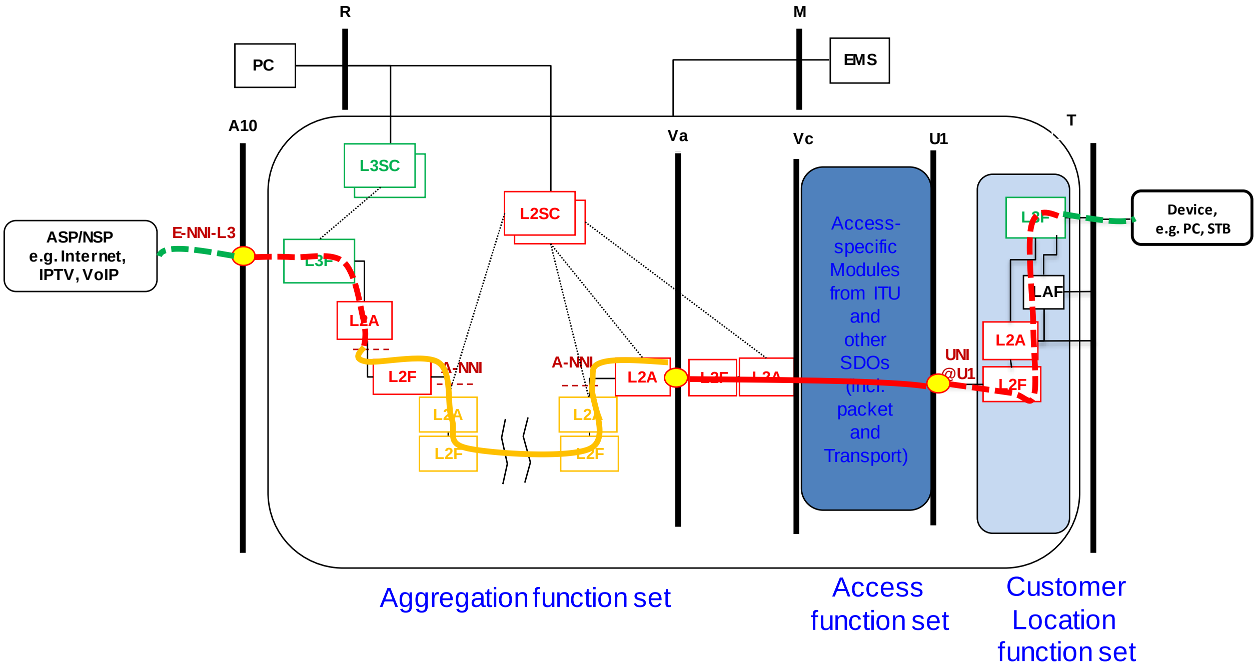

Figure 56 Residential Internet Access (centralized)

- Ethernet Service Layer connectivity between L3F at CLF and L3F across Access and Aggregation Function Set (red line between UNI@U1 and ENNI-L3@A10)
- IVC between I-NNI and I-NNI (orange line between SI-NNI and I-NNI@Va)
- IP/PPP termination in the L2A/L3@SI-NNI function with associated session control (AAA)

Option 2

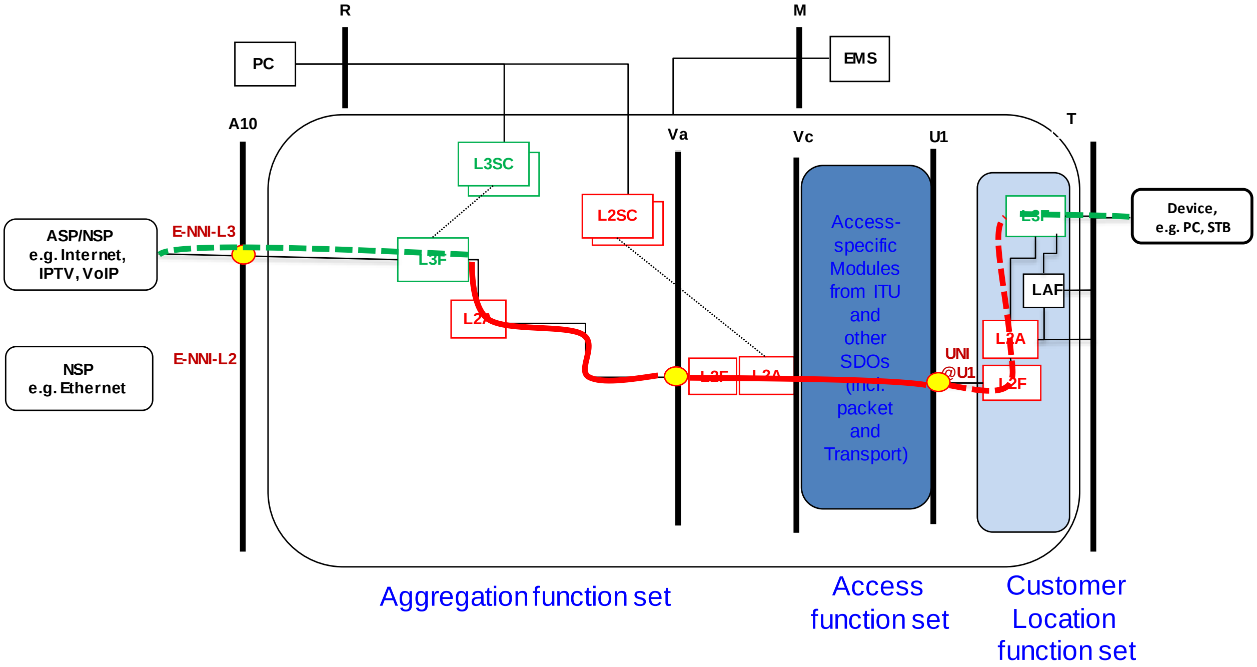

Figure 57 Residential Internet Access (Distributed)

- Native Ethernet in access (L2A does S-VLAN imposition).
- Null IVC , in other words there is no Supporting Aggregation Layer for that Service
- IP/PPP termination in L2A/L3F function with associated session control (AAA)

II.8.2. Residential IPTV/VoD and/or Voice

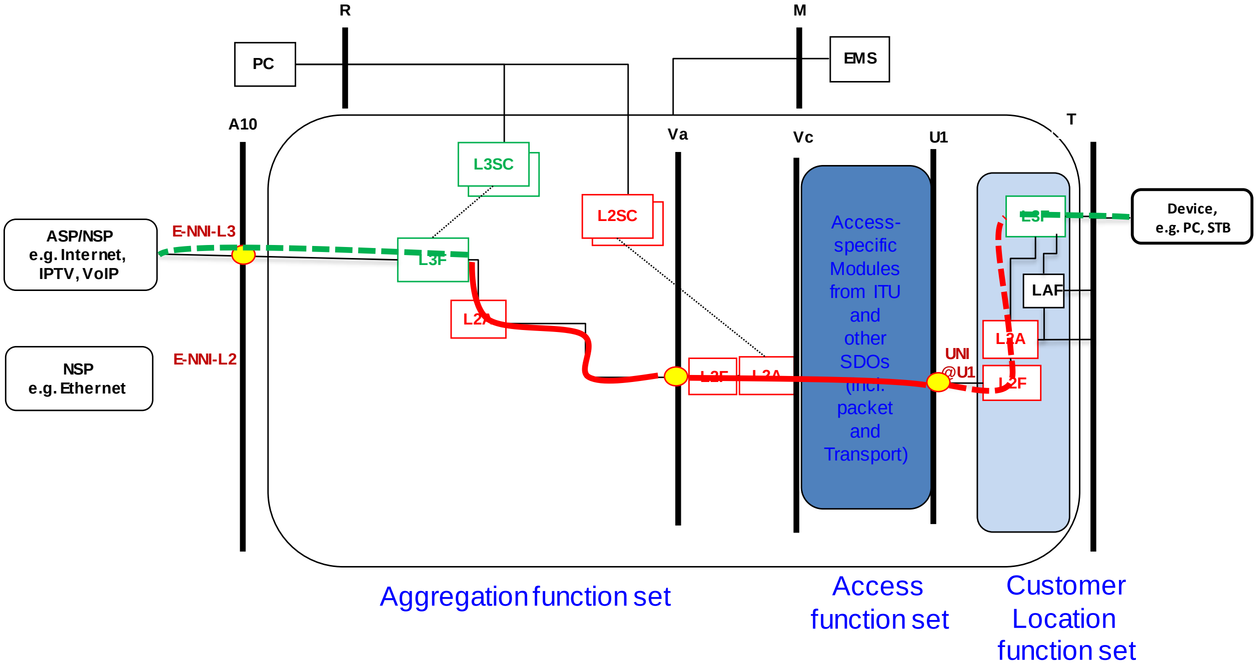

Figure 58 IPTV/VoD/Voice Service: Distributed

- Null IVC → IP Aggregation
- No need for explicit per subscriber session control
- STF: CAC, Per Service QoS (DiffServ)

Note: This use-case is independent of whether IPv4 or IPv6 is used. Also L2F, L2A and L3F can all be instantiated in the same physical node if need be.

II.9. Topology use cases

This section provides various topology use cases driven by higher bandwidth and high resiliency. Note, the use cases presented here omit L2SC/L3SC functional modules in the figures for simplicity, except when those functions are used specifically for the topology aspects.

Ethernet Service Layer Topology

Motivation:

- In order to meet the demand for increasing bandwidth over time, ANs supporting Copper access media will tend to be deployed closer to the subscriber premises.

Description:

- ANs located deeper in the network will lead to more “Tree” based distributed access network, e.g. Point to multi-point like PON – or daisy chained individual ANs

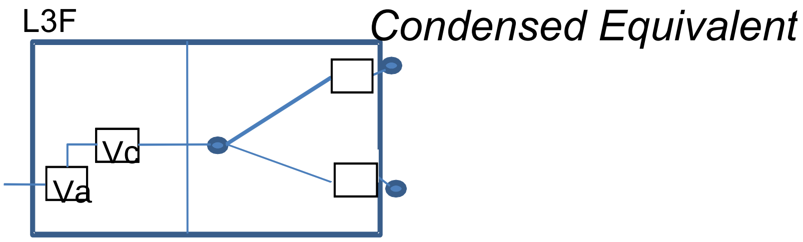

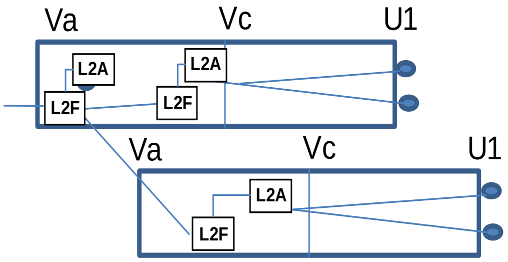

Daisy-chained ANs

  – Distributed Access Nodes may have integrated, subtended remote shelves, line cards or sealed modules.

- The logical connectivity supported is:
    1. Point to Point: e.g. E-Line

2. (Rooted) Point to Multipoint e.g. E-Tree
3. Multipoint to Multipoint e.g. E-LAN
4. Multi-layer topologies (L1, L2, L3), where L1 is emulated over a packet L2/L3 topology.

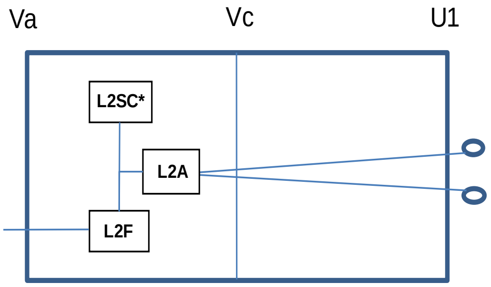

(3) LAN

L2SC* Used for authenticating access

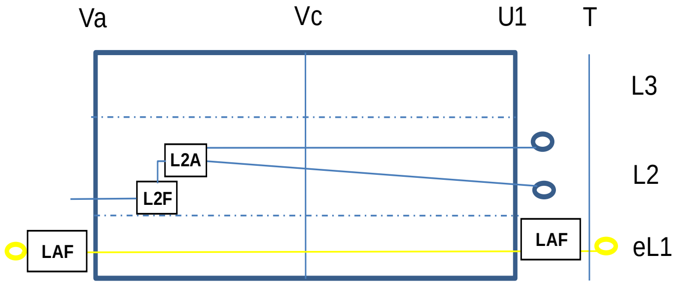

(4) Multiple Layer

- Redundancy Options:
  - Rings of access nodes (to improve overall availability with minimal extra trunking)
  - Meshing (Dual homing to the same or a different AN – the latter is primarily to support business customers who need full redundancy)

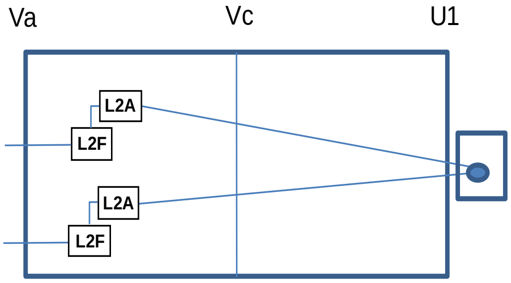

Mesh (Dual homing)

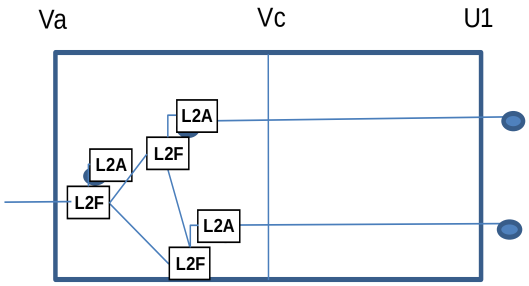

Ring type AN network

- Chains of access nodes connected to 2nd stage Aggregation or to redundant services edges to improve overall availability with minimal extra trunking

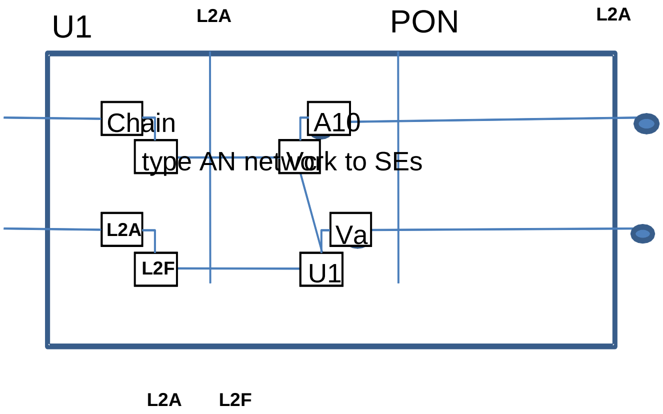

Supporting Aggregation Layer

All the physical topologies described in the access network section also apply to the logical topologies that the Supporting Aggregation layer needs to support.

In addition logical multi-rooted E-TREE topologies need to be supported e.g. for multiple service edges . Note that the logical topologies may be layer dependent.

Driven by **redundancy** considerations, the Supporting Aggregation Layer may also need to support:

- Dual homing of Access Nodes. This requires additional control and signaling for load sharing and/or switch-over, and this applies to both dual-homed device and dual-homed network configurations.
- Redundant Service Edges including scenarios where they are not directly connected to the Access Node. This will require control/signaling for switchover and fast re-route.

From a **logical connectivity** perspective, the Supporting Aggregation Layer needs to support the same characteristics as listed in the Ethernet Service Layer section but also the appropriate mechanisms to achieve connectivity to redundant Service Edges.

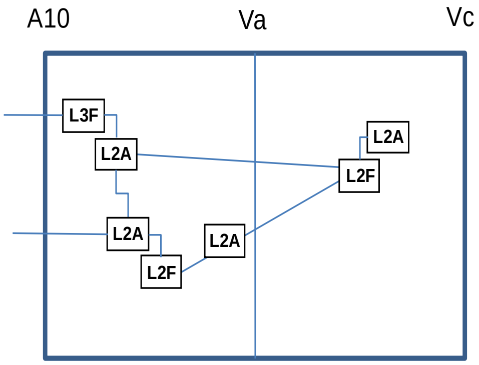

Reverse Tree (to N SEs)

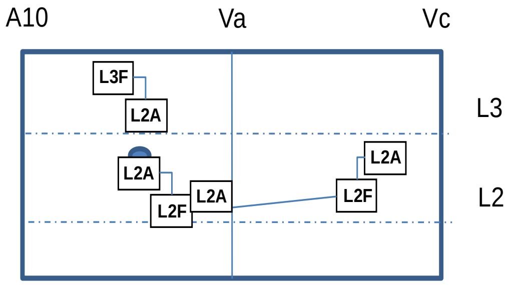

Multiple Layer SEs

- Logical connectivity
  - Multipoint to multipoint (e.g. E-LANs)

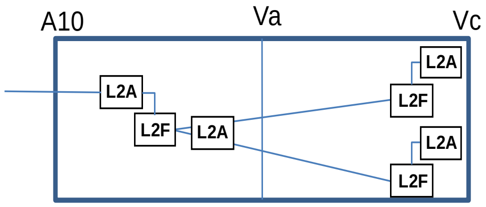

Mpt-mpt (E-LAN)

  - Dual-homed E-Trees (on multiple edges or on the same, redundant service edge)

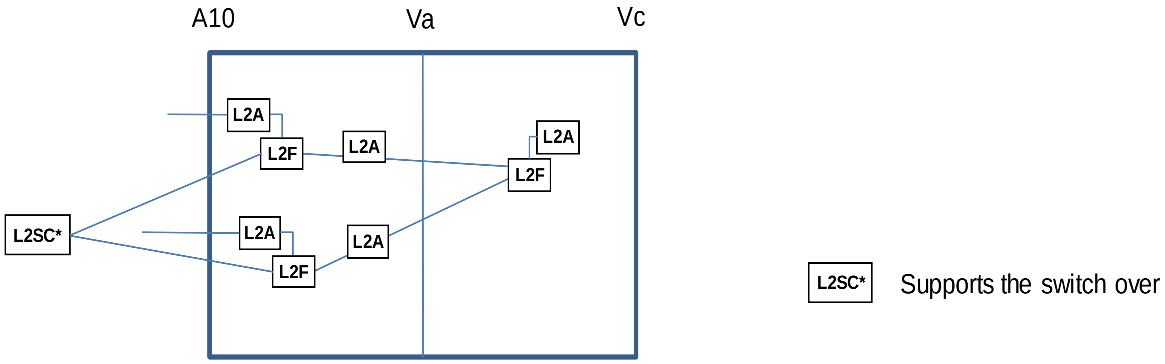

Redundancy Options:

- Dual homing at the A-NNI. This requires additional control and signaling for load sharing and/or switch-over.

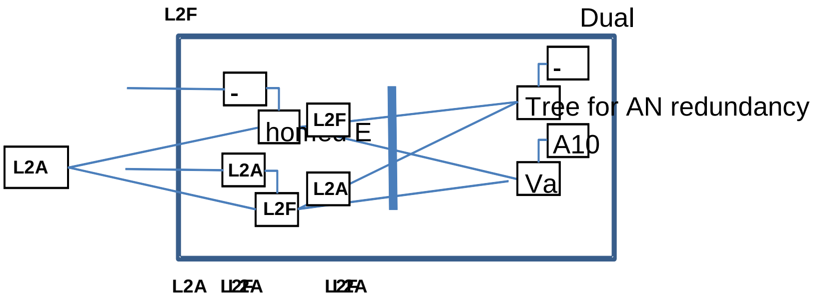

- Mesh or dual homing (diverse logical routes in the Supporting Aggregation Layer to various service edges) for the business/consumer is desired. This will require control/signaling for switchover and fast re-route.

Appendix III - Supporting users with duplicate MAC addresses

III.1. Scope

The requirement to provide service to users with duplicate MAC addresses was first identified in TR-101 (R-89). It was repeated for TR-156 (R-114) and for TR-145. This appendix describes a method for supporting this requirement, i.e. for ensuring that multiple end-users having customer premises equipment with identical MAC addresses are not denied service access on that account.

The method is denoted MAC Address Translation (MAT). With this method each end-user MAC address exposed on the U-interface is symmetrically replaced by a unique administered MAC address at the V-interface. This ensures uniqueness of MAC addresses at the V-interface, while identical MAC addresses may be used at different U-interfaces.

III.2. Background

An Ethernet-based access and aggregation network may rely on end-user MAC addresses for forwarding downstream traffic. This is typically the case in network scenarios deploying an N:1 VLAN architecture in which multiple end-users share a single broadcast domain (VLAN). Simplified subnet and VLAN administration, and well as ability to easily support multi-edge architectures, are some of the motivations for deploying N:1 VLAN architectures. Here, the end-user MAC address is used by aggregation switches and access nodes to determine on which egress port to forward the corresponding Ethernet frame. However, this forwarding method requires all MAC addresses in use to be unique within each broadcast domain. If multiple end-users are simultaneously using the same MAC address, the downstream forwarding behavior becomes unpredictable, network and service access may be interrupted, and confidentiality of end-user traffic is jeopardized.

Although the IEEE MAC address concept is originally designed to ensure global MAC address uniqueness, multiple end-users deploying identical MAC addresses may occur, either as the result of experimental or malicious end-user behavior, or by negligence in the CPE manufacturing process. State-of-the-art CPE routers offer the ability to easily modify the MAC address of the device. For a bridged CPE scenario, many terminal types such as PCs offer a similar option for their NICs. In other scenarios, CPE routers or NICs may from the factory be configured with identical MAC addresses. In conclusion, end-user MAC addresses can never be trusted as addressing points. This aspect must be dealt with whenever end-users share a broadcast domain, for example N:1 VLAN access architectures.

III.3. MAT method description

The MAT method establishes a one-to-one relation between on one side the MAC addresses of end-user devices exposed on each U-interface (1st mile), and on the other side the MAC addresses used for device identification and traffic forwarding in the Ethernet aggregation network. The translation between MAC addresses is performed by the Access Node, i.e. at the border between access lines used by single customers and the aggregation network shared between multiple customers, thereby preventing end-users MAC addresses to be exposed on the V-interface.. Only unicast addresses are to be mapped; multicast and broadcast addresses are never used as source MAC addresses and therefore not an object of spoofing.

In the picture below, an end-user device sends an Ethernet frame towards the BNG with MAC address B (destination address of the frame). The end-user device has MAC address U1 (source address of the frame). The Access Node converts the source address U1 to its alias address V1. In the opposite direction, the Access Node converts the destination address of a downstream frame from V1 to U1.

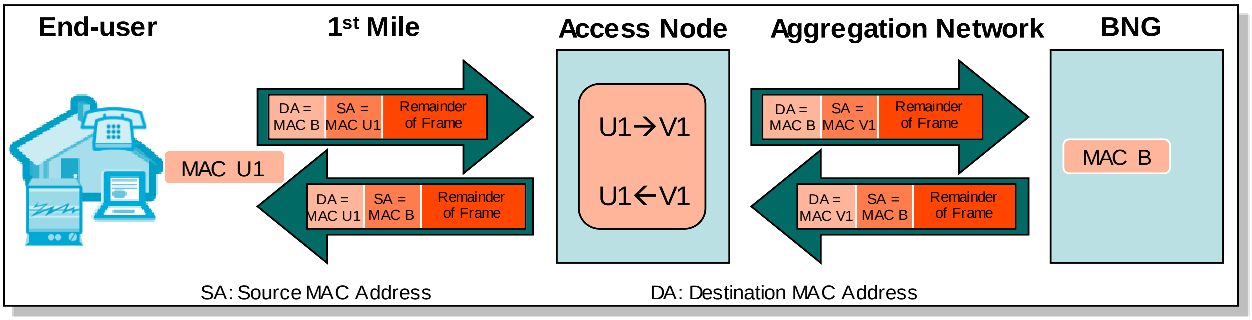

Figure 59 MAC Address Translation

III.4. VMAC Address Layout

An alias MAC address is also denoted a Virtual MAC address, or just VMAC. The format of a MAC address is depicted below.

Following the IEEE 802 address format, the first byte of the address includes a 1-bit field indicating if the address is globally administered (i.e. must be globally unique), or if the address is locally administered. In the latter case the value of the address field (46 bit) can be chosen freely by the administrator of the Ethernet broadcast domain. Address uniqueness must, however, still be guaranteed within each broadcast domain of the aggregation network.

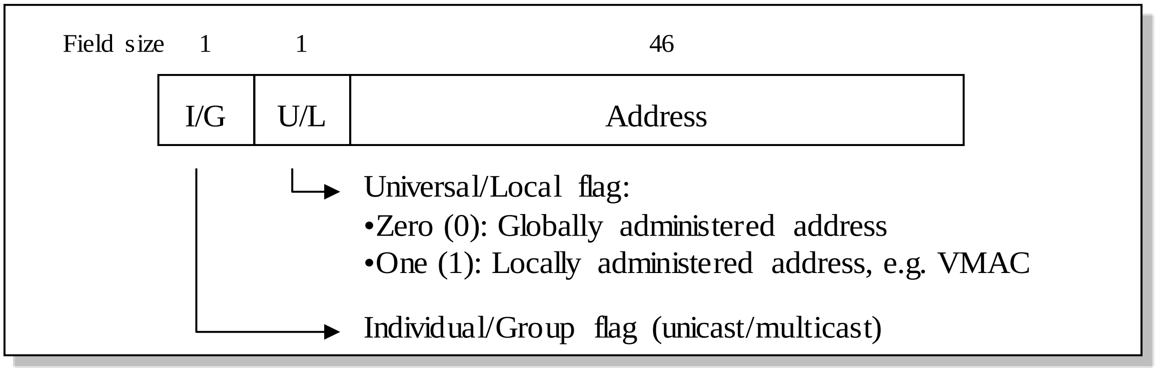
Field size
1
1
46
I/G
U/L
Address
Universal/Local flag:
•Zero (0): Globally administered address
•One (1): Locally administered address, e.g. VMAC
Individual/Group flag (unicast/multicast)

Figure 60 IEEE 802 address format

The address uniqueness requirement and a general desire to avoid manual configuration both call for defining an address assignment algorithm that each Access Node will use autonomously in the generation and assignment of VMACs.

To ensure the VMACs generated by different Access Nodes operating in the same broadcast domain do not overlap, VMAC generation algorithms need to process unique Access Node Identifiers, which must be configurable by the operator and may be extracted from the Access Node general configuration (cf. TR-101 Table 2 Access_Node_ID in the Circuit ID Syntax). User port numbers may also be used by VMAC generation algorithms, allowing a user port to be directly identified from a VMAC.

Multiple algorithms can be designed to ensure that these requirements are met, and the optimal layout for a given network depends on a number of parameters, including number and maximum size (number of end-user ports) of Access Nodes in the network, and the number of MAC addresses allowed per end-user. It is beyond the scope of this document to settle for one specific algorithm.

III.5. Impact on protocols using End-user MAC addresses in payload

Some protocols use the end-user MAC address within the payload of the Ethernet frame, for example because the address value carries relevant information to be processed by the frame receiver at a protocol layer higher than the MAC layer.

III.5.1. ARP

The ARP protocol comprises a field for the host MAC address: Sender Hardware Address (SHA). This is used both in ARP request and reply. Thus, for an upstream ARP message the Access Node must convert the SHA to the associated VMAC.

III.5.2. DHCP

The DHCP protocol comprises a field for the host MAC address: Client Hardware Address (CHADDR) which is contained in all DHCP messages. The client inserts this address in the payload of its Discovery messages and Request messages, and the DHCP server/relay uses the given CHADDR as destination MAC address in its replies (Offer, Acknowledge). Consequently, in order for the DHCP reply to reach its destination, the Access Node should change the original CHADDR value to the VMAC address. If the DHCP server requires knowledge about the original CHADDR for other purposes, e.g. device identification, that value can be carried in DHCP option 61 (DHCP Client Identifier), inserted by the Access Node.

III.5.3. PPPoE

Using MAT in PPPoE scenarios has no specific impact on the Access Node functionality. The original client MAC address is considered to be not relevant to the BNG. Alternatively, the MAC address could be conveyed in the relay-session-id of the PADI message.

III.5.4. CFM

The CFM protocol (IEEE802.1ag) is impacted by MAT. More specifically, the LTM message includes the MAC Address in the payload, enabling LTRs back to the originating MEP. Here, the Access Node must modify the payload with the VMAC address.

It also has to be noted that MAT prevents testing end to end Ethernet continuity in the full path: a failure in the forwarding path could potentially be undetected by link-trace CFM, in case the AN would have stopped MAT functions for basic Ethernet frames but not for CFM messages. Note this is similar to the limitations associated with using an interworking function between Ethernet OAM and ATM OAM for CPE not supporting Ethernet OAM (TR-101 §7.1).

III.5.5. IPv6

MAT used with IPv6 over Ethernet impacts the following control messages of ICMPv6 (RFC4443, RFC4861, RFC3122, RFC2710, RFC3810):

- Neighbor solicitation, Neighbor advertisement (MAC address of sender included in payload)
- Router solicitation, Router advertisement (MAC address of sender included in payload)
- Inverse ND solicitation, Inverse ND advertisement

For these protocols, the MAC address of the sender/receiver appears inside the Payload and must be dealt with similar to ARP and DHCP for IPv4.

DHCPv6 also comprises an option for the client’s MAC address. As with DHCPv4 this option must be mapped to the VMAC address by the Access Node.

Since MAT changes a possible relation between MAC and Link Local addresses, a BNG must not assume any such relation (i.e. LLA are not derived from VMAC).

III.6. End-user Identification

Due to the structured layout of the VMAC address, each upstream Ethernet frame carries information about the end-user port on which it was received, and each downstream Ethernet frame carries information about the end-user port on which is sent. This information may be used as unique end-user port identifier.

As it is replaced by the VMAC, the CPE MAC address is not any more present as source/destination of upstream/downstream Ethernet frames.

In other words, with MAT any Ethernet frame can be uniquely and directly related to a specific end-user port; as opposed to without MAT, where any Ethernet frame can be directly related to a specific end-user device.

III.7. MAT Deployment Considerations

MAT solves the MAC uniqueness requirement for N:1 VLAN scenarios where MAC learning is fundamental to traffic forwarding. MAT supports deployment cases where a VLAN is shared between several customers, and also cases where a VLAN is shared between several Access Nodes.

MAT supports cases where IP antispoofing is performed in the BNG, as well as cases where IP antispoofing is performed in the Access Node.

As opposed to rejecting service to devices with a “duplicate MAC address”, MAT enables any device MAC address to be used, and avoids considering as duplicates MAC addresses moving between user ports within the same N:1 VLAN.

MAT supports deployment cases where the RG is bridged, as well as cases where the RG is routed.

MAT supports deployment cases where the RG is provided by the customer and purchased from the retail market, as well as cases where the RG is provided by the operator.

III.8. Administration and Management Considerations

As MAT allows providing service to users with duplicate MAC addresses, it may reduce hotline and administrative overheads by avoiding some situations where customers would be denied service.

In a 1:1 VLAN architecture, as in TLS architecture, MAC uniqueness is normally not a problem, and MAT is therefore not required in these architectures. Likewise MAT may be deployed for some N:1 VLAN where needed, while it may not be needed in some other N:1 VLAN, e.g. multicast. These VLAN architectures, i.e. 1:1 VLAN, TLS, N:1 VLAN with MAT and N:1 VLAN without MAT, can coexist side by side. The access node must support activating and deactivating MAT per VLAN.

For troubleshooting purposes, the translation table must be readable per user port and VLAN on the Access Node. MAT does not have any impact on testing and troubleshooting procedures, as long as the CPE’s MAC is not used in these procedures, in which case it would require using the VMAC instead.

The Access Node MAT functions may have to be updated when new protocols using CPE’s MAC addresses in their payload are deployed.

A unique VMAC pool per AN needs to be provisioned, i.e. a prefix, unique per broadcast domain, that will be used by the VMAC generation algorithm. This prefix may be automatically derived from the Access Node Identifier already provisioned for generating the DHCP Circuit-ID.

End of Broadband Forum Technical Report TR-145