TECHNICAL REPORT

TR-345
Broadband Network Gateway and Network Function Virtualization

Issue: 1
Issue Date: October 2016

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Technical Report has been approved by members of the Forum. This Technical Report is subject to change. This Technical Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Technical Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Technical Report, or use of any software code normatively referenced in this Technical Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Technical Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Technical Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Technical Report except to the extent this Technical Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	17 October 2016	14 December 2016	Oliver Thorp, Sky	Original

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

Editor	Oliver Thorp	Sky	oliver.thorp@sky.uk
Architecture and Migration Work Area Directors	Dave Thorne Dave Allan	BT Ericsson	david.j.thorne@bt.com david.i.allan@ericsson.com

TABLE OF CONTENTS

List of Figures

List of Tables

Executive Summary

Network operators are deploying Virtual Network Functions (VNFs) that either augment or replace the Multi-Service Broadband Network Gateway (MS-BNG) in deployed TR-101 [1] or TR-178 [5] network architectures. This Technical Report describes how VNFs and their supporting Network Function Virtualization Infrastructure (NFVI) can be integrated with these Broadband Forum architectures. This includes scenarios where NFVI is connected directly to a TR-101 access network and also where VNFs are deployed behind an MS-BNG as part of a service graph.

The use cases considered include Ethernet, IP and tunneled presentation of traffic into NFVI. The Technical Report addresses support for a Network Enhanced Residential Gateway (TR-317) [10], Virtual Business Gateway (WT-328) [11] and Virtual L2TP Network Gateway. However, the TR-345 framework should be applicable to any other Broadband Forum architecture using tunneled connectivity to a service edge.

The main new features in this Technical Report are traffic steering to virtual functions, the introduction of NFVI Gateways and requirements for VNFs integrating with TR-101 and TR-178 architectures.

1 Purpose and Scope

1.1 Purpose

The purpose of this Technical Report is to define a framework for the introduction of Network Function Virtualization into TR-178 architectures. This is motivated by NFV being used to:

- Add new service functions to existing MS-BNG platforms
- Provide greater flexibility with regard to where the subscriber session is terminated to existing MS-BNGs
- Replace the physical MS-BNG entirely.

The introduction of NFV into Broadband Forum specifications provides a potential opportunity for a complete re-architecting of the Multi-Service Broadband Network. However, it will be possible to gain some benefits from NFV deployment without changing out the large number of existing access and aggregation network elements. The introduction of virtualization should also not require a complete reengineering of existing processes and procedures. Therefore, as a migration step, this Technical Report provides options for incremental deployment by seeking to align the capabilities and operational practices with currently deployed broadband regional access networks.

1.2 Scope

This Technical Report augments TR-178 to provide migration paths for the deployment of NFV infrastructure in existing BBF networks.

This Technical Report specifies the function of an NFVI Gateway in Broadband Forum architectures. The role of the NFVI Gateway is to provide interworking between the TR-178 aggregation network and the NFVI.

Examining the implications of introducing virtualization into the TR-178 architecture with regard to:

- Extending the TR-101/178 Ethernet service layer into NFV infrastructure
- Service chaining of selected flows from a subscriber IP session into the first VNF in the NFV infrastructure
- Traffic management and QoS
- Load Sharing

The following items are out of scope:

- VNF definition
- Service & Cloud management and orchestration
- Virtual Networks within the NFVI Network Domain
- NBI/SBI definitions and protocols (e.g. SDN)
- Cloud Central Office (including Virtualized Access Node)

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 [17].

MUST	This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase 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.

The requirements have been marked to show if the requirement should be provisioned or managed using a traditional management interface, or by a control protocol or both. **M** is used to denote that it is provisioned by a Management Plane, while **C** is used to denote that is configured by a Control Plane.

2.2 References

The following references are of relevance to this Technical Report. At the time of 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-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[2]	TR-145	*Multi-service Broadband Network Functional Modules and Architecture*	BBF	2012
[3]	TR-146	*Subscriber Sessions*	BBF	2013
[4]	TR-177	*IPv6 in the context of TR-101*	BBF	May 2010
[5]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[6]	TR-187	*IPv6 for PPP Broadband Access*	BBF	November 2010
[7]	TR-221	*Technical Specification for MPLS in Mobile Backhaul Networks*	BBF	October 2011
[8]	TR-224	*Technical Specification for MPLS in Carrier Ethernet Networks*	BBF	September 2014
[9]	TR-242 Issue 2	*IPv6 Transition Mechanisms for Broadband Networks*	BBF	February 2015
[10]	TR-317	*Network Enhanced Residential Gateway*	BBF	July 2016
[11]	WT-328	*Virtual Business Gateway*	BBF	WIP
[12]	ETSI GS NFV 002 V1.2.1	*Network Functions Virtualisation (NFV); Architectural Framework*	ETSI	2014
[13]	ETSI GS NFV-INF 001 V0.3.8	*Network Functions Virtualisation; Infrastructure Overview*	ETSI	March 2014
[14]	ETSI GS NFV-INF 005 V1.1.1	*Network Functions Virtualisation (NFV); Infrastructure; Network Domain*	ETSI	December 2014
[15]	ETSI GS NFV-MAN 001 V0.5.0	*Network Function Virtualization (NFV) Management and Orchestration*	ETSI	May 2014
[16]	ETSI GS NFV-	*Network Functions Virtualisation (NFV); NFV*	ETSI	October

	SEC 001 V1.1.1	*Security; Problem Statement*		2014
[17]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[18]	RFC 6085	*Address Mapping of IPv6 Multicast Packets on Ethernet*	IETF	2011
[19]	RFC 6788	*The Line-Identification Option*	IETF	2012

2.3 Definitions

The following terminology is used throughout this Technical Report.

Edge Network Function The first upstream and last downstream functional component that enforces policies associated with a Subscriber Session

ENF System The set of ENF instances and their management that serves one or more S-VLAN delineated broadcast domains in the ENF hosted in NFVI deployment scenario

I-NNI Internal Network-to-Network Interface; as defined in MEF.4

IP Flow An IP Flow is identified by a 5-tuple IP parameter traffic classifier. An IP Flow identifier forms the classification element of a traffic policy that is applied to a Subscriber Session. The 5-tuples is made up of following header fields: source IP address, source port, destination IP address, destination port and protocol. (as defined in TR-146 [3]).

IP Session A grouping of traffic according to one or more classifiers visible at a control point, called the IP Edge, in the broadband network. The classifier is composed of, at a minimum, a Subscriber's IP address (v4 or v6), IPv4 subnet or IPv6 prefix (as defined in TR-146).

MS-BNG TR-178 introduces the Multi-Service BNG (MS-BNG) which extends the capabilities of a traditional BNG to offer services to both residential and business customers as well as to allow mobile backhaul deployments. To achieve this, it performs Ethernet Aggregation and can either forward packets via MPLS or through IP Aggregation/routing. A MS-BNG is part of a TR-145 [2] network architecture and can be deployed in a hierarchical MS-BNG architecture.

L2 classifiers Layer-2 header fields used to identify and/or classify traffic for further action such as QoS enforcement or L2 forwarding policy. The Layer-2 classifiers used are:

- Source MAC address
- Destination MAC address
- 802.1Q/p markers (including C/S-VLAN when stacked VLAN are used)
- Various Ethertypes (IPv4, IPv6, PPPoE, etc).

L3 classifiers	Layer-3 header fields and some Layer-4 header fields used to identify and/or classify traffic for further action such as QoS enforcement or L3 forwarding policy. The Layer-3 classifiers used are: • Source IP address • Destination IP address • DSCP field • IP Protocol numbers (TCP or UDP) • Source Port Number (TCP or UDP source port number) • Destination Port Number (TCP or UDP destination port number)
NFVI Gateway Function	The function that interworks the MAN or WAN protocols (e.g. L2, L3 VPN instances) to Virtual Networks used internally to the NFVI
SF	Service Function, a VNF or PNF that provides a network service
SFC	Service Function Chain, a list of ordered service functions through which traffic flows are steered
Subscriber Session	A Subscriber Session is either a PPP Session, an IP Session, or an Ethernet Session. Subscriber sessions are used to represent all traffic that is associated with that subscriber by a given service provider in order to provide a context for policy enforcement (as defined in TR-146).
Va	Reference point at which the first level of Ethernet aggregation and the rest of the network interconnect. 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. distributed wholesale handoff). In TR-178, an Access Node with an internal Va reference point will use the V reference point for its uplinks.
Virtual Network	A Virtual Network is the network construct that provides network connectivity to one or more VNFs that are hosted on the NFVI.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AAA	Authentication, Authorization, Accounting
AFTR	Address Family Translation Router
AN	Access Node
BGP	Border Gateway Protocol
BNG	Broadband Network Gateway
BR	Border Router
BSG	Broadband Service Gateway
BSS	Business Support Systems
CF	Classifier Function
CPE	Customer Premises Equipment.

C-Tag	C-VLAN Tag
C-VLAN	Customer VLAN
EFP	Ethernet Flow Point (TR-145)
EMS	Element Management System
ENF	Edge Network Function
EPL	Ethernet Private Line
EVC	Ethernet Virtual Connection
EVPL	Ethernet Virtual Private Line
FCAPS	Fault, Configuration, Accounting, Performance, Security
GRE	Generic Routed Encapsulation
GTP	GPRS Tunneling Protocol
GWF	NFVI Gateway Function
IGMP	Internet Group Management Protocol
I-NNI	Internal Network Network Interface
IP	Internet Protocol
L2F	Layer 2 Forwarding
L2TP	L2 Tunneling Protocol
L2TS	L2 Tunnel Switch
LAC	L2TP Access Concentrator
LAN	Local Area Network
LNS	L2TP Network Server
MANO	Management and Orchestration
ME	Maintenance Entity
MEF	Metro Ethernet Forum
MEP	Maintenance End Point
MPLS	Multi-Protocol Label Switching
MP-TCP	Multi-Path TCP
MS-BNG	Multi Service BNG
NFV	Network Functions Virtualization
NFVI	NFV Infrastructure
NFVI-GW	NFVI Gateway
NLIM	Network Line Interface Module
OAM	Operations Administration and Maintenance
OSS	Operations Support Systems
PADI	PPPoE Active Discovery Initiation
PADO	PPPoE Active Discovery Offer

POP	Point of Presence
PPP	Point to Point Protocol
PPPoE	PPP over Ethernet
PW	Pseudo Wire
QoS	Quality of Service
RG	Residential Gateway
S-Tag	S-VLAN tag
S-VLAN	Service VLAN
TCP	Transmission Control Protocol
TR	Technical Report
TWAG	Trusted Wireless Access Gateway
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VNF	Virtual Network Function
VPN	Virtual Private Network
VIM	Virtual Infrastructure Manager
WLAN	Wireless Local Area Network
WT	Working Text

3 Technical Report Impact

3.1 Energy Efficiency

Migration of service functions to NFV infrastructure provides the opportunity to move from a service deployment model with dedicated equipment in physical MS-BNG nodes to a model where functions are provided using compute infrastructure. The compute infrastructure can be realized through NFV optimized server architectures and dynamically dimensioned based on load. This can provide the opportunity to reduce the hardware footprint for a service and reduce the energy consumption.

Network functions need to be allocated to the appropriate platform to benefit from these efficiencies. Poor mapping of forwarding intensive functions to compute optimized NFVI rather than data plane optimized resources could result in a drop in energy efficiency. The carrier will need to deploy methodologies to reduce energy consumption by leveraging the flexibility of NFV to deploy sophisticated power management such that the overall power consumption is reduced rather than increased.

3.2 IPv6

The IPv6 transition mechanisms described in TR-242 [9] (including Dual-Stack CGN, DS-Lite AFTR and 6RD BR) are candidate service functions for virtualization. Virtualization of these functions allows flexibility to scale and distribute them as required during the IPv6 transition and reduce the dependency on the physical MS-BNG space, power and processing capacity.

3.3 Security

Introduction of NFV into a service provider environment increases the attack surface that needs to be managed by the service provider. The ETSI NFV ISG has published a problem statement [16] defining the potential areas of concern that need to be considered as part of any NFV deployment. This includes generic virtualization threats and generic networks threats in addition to NFV-specific threats. Methods to secure the connection to the NFVI and the NFVI itself are out of the scope of this document.

3.4 Privacy

NFV can enable some functions of the Residential Gateway (NERG use case) or functions of the Business Gateway (vBG use case) to be moved from the customer premises to service provider hosted infrastructure. This can create privacy risks. This is discussed in related specifications, TR-317 and WT-328.

Flexible service chaining permits distribution of the service edge into multiple nodes and can increases the number of functions in the network with subscriber awareness. Each of these functions and the communication of subscriber information across the functions will need to be managed to ensure compliance with privacy regulations.

4 Overview of Fundamental Architectures and Topologies

This Technical Report defines a framework for the introduction of Network Function Virtualization into TR-101 and TR-178 based architectures.

4.1 Motivations for NFV deployment

This Technical Report considers NFV deployment models where Virtual Network Functions (VNFs) augment or replace the MS-BNG. These deployment models are described generically below, with more specific examples in section 5 and Appendix I.

4.1.1 NFV deployed to add services to existing MS-BNGs

This scenario permits the introduction of NFV Infrastructure (NFVI) into the network to deliver value added services that are not supported by the currently deployed MS-BNG. In this scenario the IP session management functions are delivered by the existing MS-BNGs and subscriber traffic is selectively steered to additional service functions deployed in the NFVI.

4.1.2 NFV deployed to add flexible subscriber session termination to existing MS-BNGs

This scenario permits the introduction of NFVI into the network in order to permit the creation of a service specific vBNG and/or support service provider trials. The MS-BNG may selectively pass the traffic of a subset of subscribers to the NFVI at layer-2. The IP session termination and any value added services for that set of subscribers are implemented in the NFVI.

4.1.3 Virtual Edge MS-BNG

This scenario supports a provider that wishes to deploy a Virtual Edge MS-BNG. An NFVI Gateway is directly connected to a TR-101 access network and the IP session / MS-BNG functional elements are implemented as VNFs.

4.2 Reference Model

Figure 1 shows the deployment reference model showing how NFV Infrastructure is integrated with the existing BBF components. The NFVI replaces the MS-BNG, is deployed beside the MS-BNG or deeper in the IP/MPLS based aggregation network. Traffic from the MS-BNG could either go through NFVI or be handed over to A10 directly.

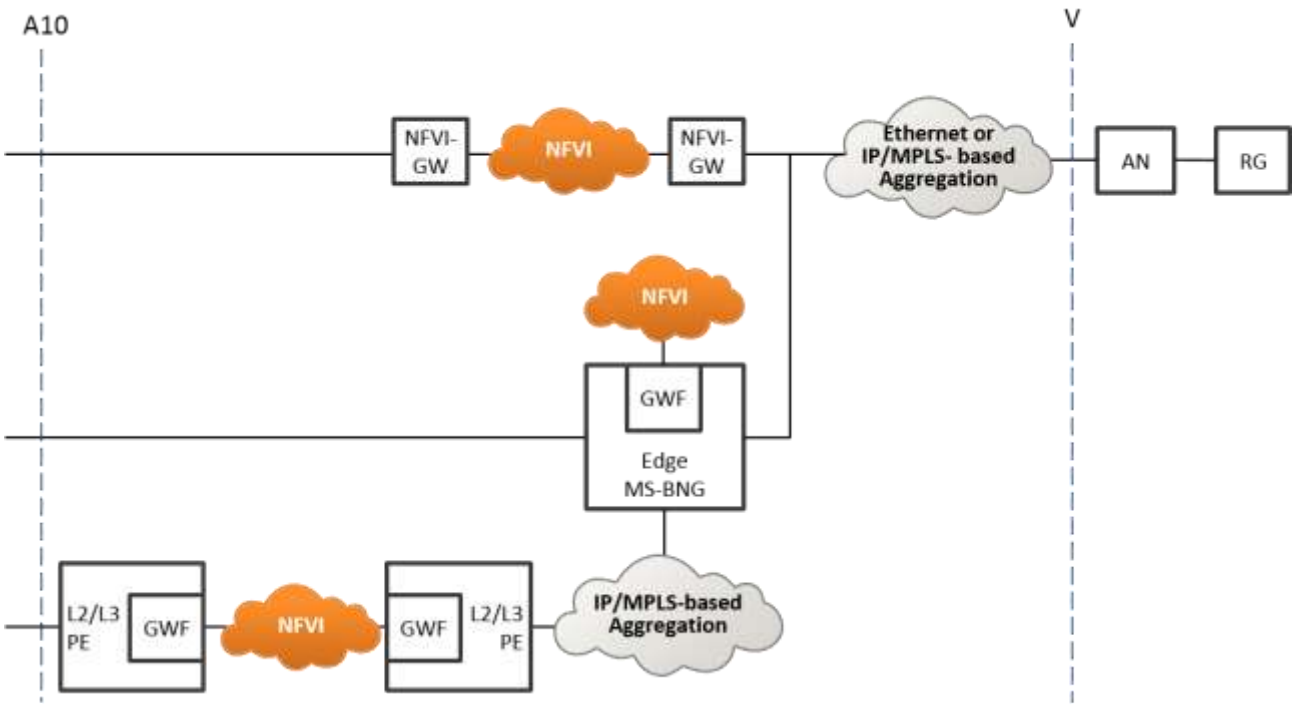

Figure 1 – NFVI deployment reference model

NFVI is deployed at a segment of the network between the A10 and V reference point where the TR-178 architecture permits an MS-BNG to reside. The virtual Edge MS-BNG, nested MS-BNG (L2 handoff) and BSG (L3 handoff) can be directly implemented as VNF(s) within the NFVI permitting incremental introduction of NFV into provider networks.

The deployment of NFVI requires a function at the boundary between the MAN or WAN and the NFVI. This function is termed the NFVI Gateway Function (GWF). Its main purpose is to interwork the MAN or WAN protocols (e.g. L2, L3 VPN instances) to Virtual Networks used internally to the NFVI. The GWF may be implemented as a stand-alone NFVI Gateway (NFVI-GW) or embedded in an MS-BNG or MPLS PE router. The NFVI Gateway Function for upstream and downstream traffic flows may be physically separated or co-located (implemented by the same network entity).

Future Broadband Forum projects may consider additional deployment scenarios for NFV in the access network and customer premises but these are out of the scope of this technical report.

Figure 2 shows the relationship between the Broadband Aggregation Network, Edge MS-BNG, NFVI Gateway Function and some of the relevant entities and interfaces in the ETSI NFV architecture [13][14][15].

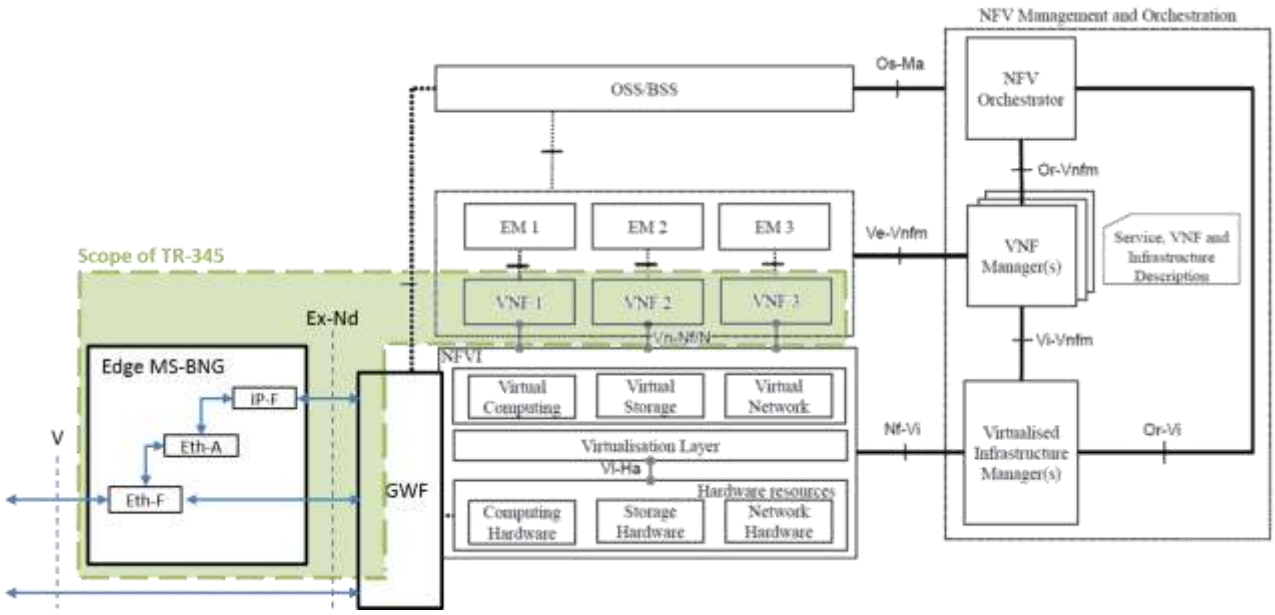

Figure 2 – TR-345 mapping onto ETSI NFV architecture

The NFVI encompasses all the hardware (e.g. compute, storage, and networking) and software (e.g. hypervisors) components, constituting 3 domains defined by the ETSI NFV ISG: Hypervisor Domain, Compute Domain and Network Domain. These together provide the infrastructure resources where VNFs are deployed. The NFVI and the NFV Orchestration and Management are out of scope of this TR. This TR focuses on the aspects of NFVI Network Domain and its interaction with the MS-BNG and/or TR-101 access networks.

The ETSI NFV architecture describes the OSS/BSS, the Element manager(s), the Virtual Network Function(s) (VNFs) and the services provided by the VNF Manager to the VNFs. The VNF Manager handles most aspects of virtualization management for the set of VNFs implemented in the NFVI. This includes fault recovery, scale out/in, lifecycle management etc. The VNF is considered to be a software image for a network function e.g. a nested MS-BNG or a BSG. A VNF instance is a copy of the software image deployed as a virtual machine on a server. The Element Manager performs FCAPS for a given VNF. These are not discussed further in this Technical Report.

The interfaces addressed by requirements in this Technical Report are:

- Ex-Nd – This reference point corresponds to the external interface to the NFVI. This is out of scope of ETSI ISG NFV specifications.

- Vn-Nf/N – The reference point between the NFVI Network Domain and the Guest Stack in the VNF instance. Network Domain services are delivered to VNFs over this reference point. The services include either IP forwarding (L3 handoff, i.e. IP session) or Ethernet Services (L2 handoff) specified in [14] correspondingly.

4.3 Service Model

The addition of NFV to the broadband network architecture creates choices about where to implement service functions previously implemented by an MS-BNG. TR-178 enables a hierarchy of MS-BNGs, where an Edge MS-BNG steers specific types of Subscriber Sessions to a centralized MS-BNG. The Edge MS-BNG performs aggregation of Subscriber Sessions and can either forward packets towards nested service edge platforms via MPLS or through IP Aggregation / routing. Other use cases are supported by IP tunnels and are summarized in Appendix I.

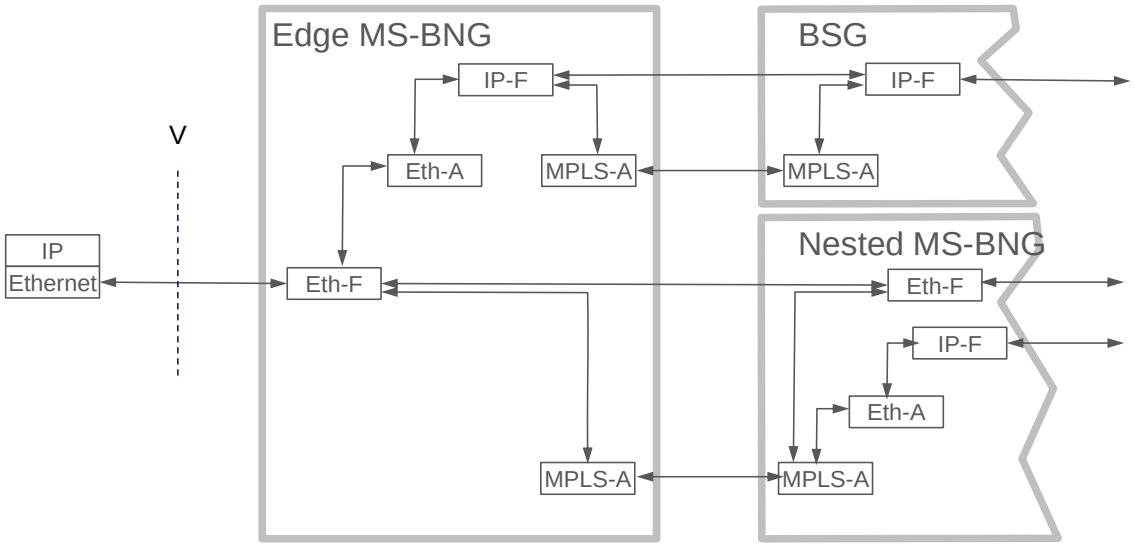

Figure 3 – TR-178 Protocol layering and high level functional distribution between Edge MS-BNG and nested service edge platforms

Figure 3 illustrates the protocol layering and uses the platform naming conventions from TR-178. This shows the possibility of a distributed implementation of an IP session by dividing functions between the Edge MS-BNG and the BSG. It also shows the Edge MS-BNG performing L2 backhaul to a more centrally deployed MS-BNG (MS-BNG Hierarchy). The diagram makes no representation as to the geographic distribution of functionality in the network or to ownership of specific platforms.

TR-345 adds the capability for the functions of these service edge platforms to be implemented using NFV. Where NFV is used to add services to existing MS-BNGs this is similar to the L3 forwarding model in TR-178. Where NFV is used to add flexible subscriber session termination, this is similar to the L2 forwarding model in TR-178.

In the case of a virtual Edge MS-BNG, some of the traditional MS-BNG functions (e.g. subscriber session termination, hierarchical traffic management) can be implemented as Virtual Network Functions.

Service Chaining is a model where traffic is forwarded selectively through ordered sets of service functions called service graphs. Examples of Service Functions include firewall, URL filtering, web-proxy, NAT, application awareness (DPI), lawful intercept, etc… These services can be applied to a subscriber, a device (see NERG use case), a flow or an application.

While these types of services have been historically delivered by either MS-BNG's or dedicated appliances, NFV helps leverage a unified infrastructure to host a variety of virtual applications. In this case, Service Functions are delivered by Virtual Network Functions.

This Technical Report describes how Broadband traffic is classified and forwarded to the first VNF in a service graph. It does not consider the protocols used to build Virtual Networks within the NFVI. This is discussed in ETSI NFV INF 005 [14].

Where a physical MS-BNG is deployed, a Classifier Function (CF) in the MS-BNG is responsible for mapping subscriber traffic to a service graph. The mapping can be applied to all subscriber traffic or specific flows. The Classifier Function can take advantage of the MS-BNG AAA/Policy server interaction to dynamically retrieve steering rules per Subscriber Session. Specifics on the service chain provisioning and protocols are out of scope for TR-345.

Figure 4 below provides an example of the TR-345 architecture where NFVI is used to support service chaining behind a physical MS-BNG. The Classifier Function within the MS-BNG is controlling how the flows within an IP session instance are forwarded to a service graph.

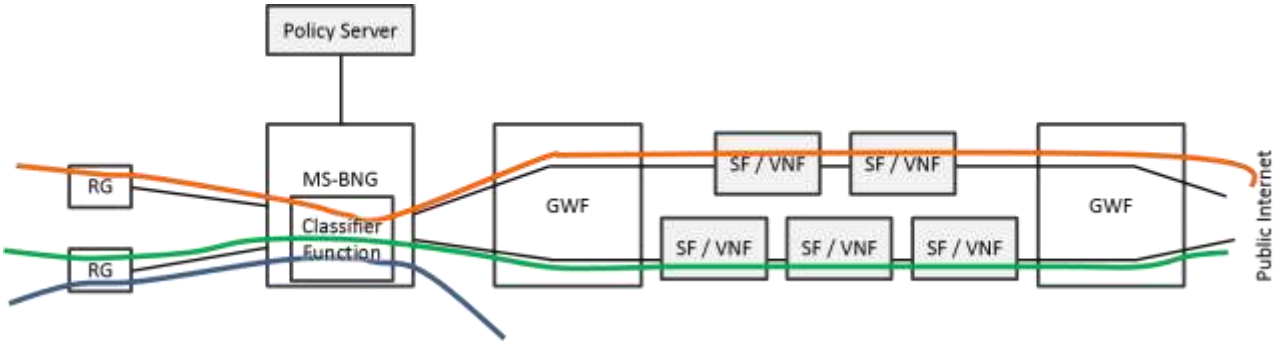

Figure 4 – Example of NFVI based service chaining

4.4 Subscriber Session Mapping

In TR-101 and TR-178, the Service Edge is a node implementing a Policy Enforcement Point for a Subscriber Session (see Figure 2 / TR-146). The implementation of a Subscriber Session may be distributed across a number of physical or virtual components.

This Technical Report introduces the term Edge Network Function (ENF) to describe part of the function performed by a Service Edge node. The ENF is the first upstream and last downstream functional component that enforces policies associated with a Subscriber Session.

The subscriber session attachment point is the point where access connectivity terminates and subscriber session processing begins. The ENF associated with a subscriber session attachment point can implement policy enforcement functions (e.g. QoS) requiring visibility of all traffic associated with the session, while subsequent VNFs may only process a subset of the traffic.

This leads to two generalized ENF deployment models: one in which the set of subscriber session attachment points is distributed inside the NFVI, and one in which they remain in a physical MS-BNG and the NFVI hosts additional services.

4.4.1 ENF and GWF Placement

An ENF can be configured in an MS-BNG (if present) or in the NFVI. The MS-BNG can initiate requests to a policy server to determine the Subscriber Session Polices that are associated with a subscriber session. These policies include whether to terminate the subscriber session locally or forward the session to another MS-BNG or an ENF hosted in NFVI. If the subscriber session is forwarded to an ENF hosted in NFVI, the ENF is implemented as part of the first VNF in a service chain.

The GWF can be embedded in an MS-BNG or it can be in an L2PE or L3PE. Figure 5 shows an example where traffic is backhauled to a centralized MS-BNG, with an ENF function on the MS-BNG and the GWF implemented on a stand-alone L2PE / L3PE. While the figure shows a single node hosting a GWF, the GWF for upstream and downstream traffic flows may be physically separated or co-located (implemented by the same network entity).

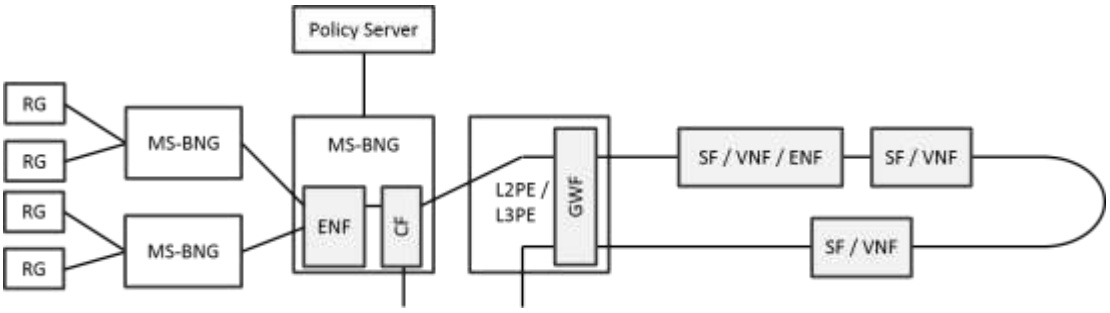

Figure 5 – Separate GWF and MS-BNG

An integrated deployment model could implement the GWF embedded in an MS-BNG. This may benefit scenarios where the NFVI is located in POPs where MS-BNGs are present. This also enables the MS-BNG to be the single exit point toward the IP backbone both for service chained and non-service chained traffic. This simplifies downstream routing.

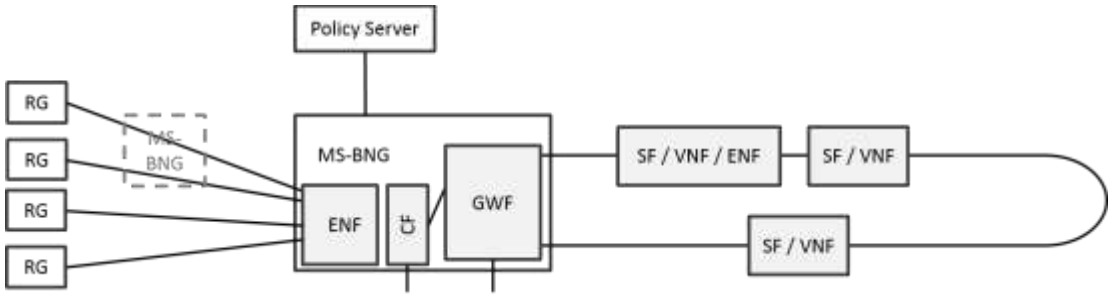

Figure 6 – GWF embedded in an MS-BNG

A final model where the NFVI Gateway is directly connected to the access network means that the ENF must be implemented as part of the first VNF in the service graph.

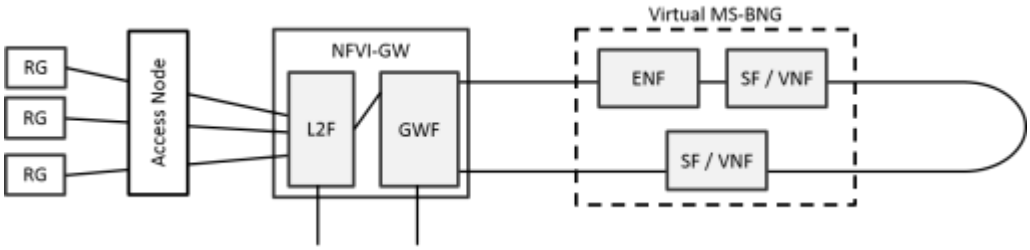

Figure 7 – NFVI Gateway directly connected to Access Network

4.4.2 ENF Hosted in NFVI

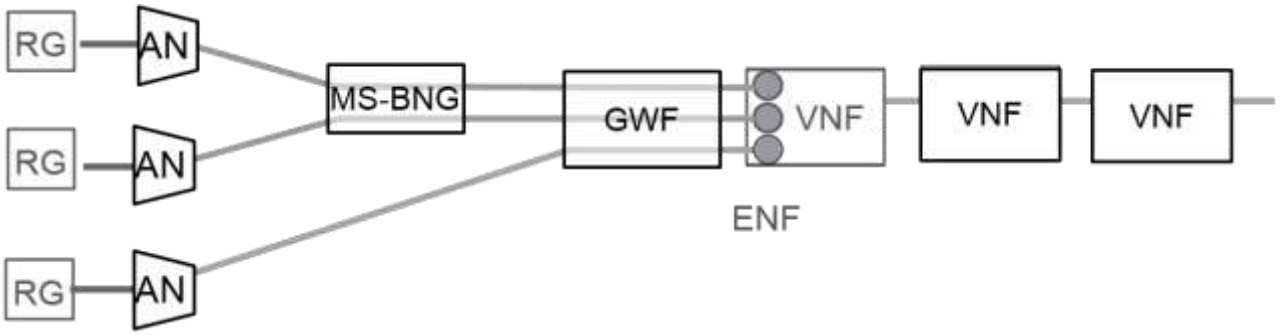

Figure 8 – ENF Model 1: ENF is hosted in NFVI

In the scenario where an ENF for the subscriber session is hosted in the NFVI, subscriber Ethernet frames are presented to the ENF across the Vn-Nf interface. This may involve some additional encapsulation, but this is beyond the scope of this document.

The Ex-Nd interface to the NFVI will include any tagging information added by the AN as well as any WAN transport encapsulation added by the MS-BNG. The ENF hosted in NFVI scenario places no new requirements on the MS-BNG.

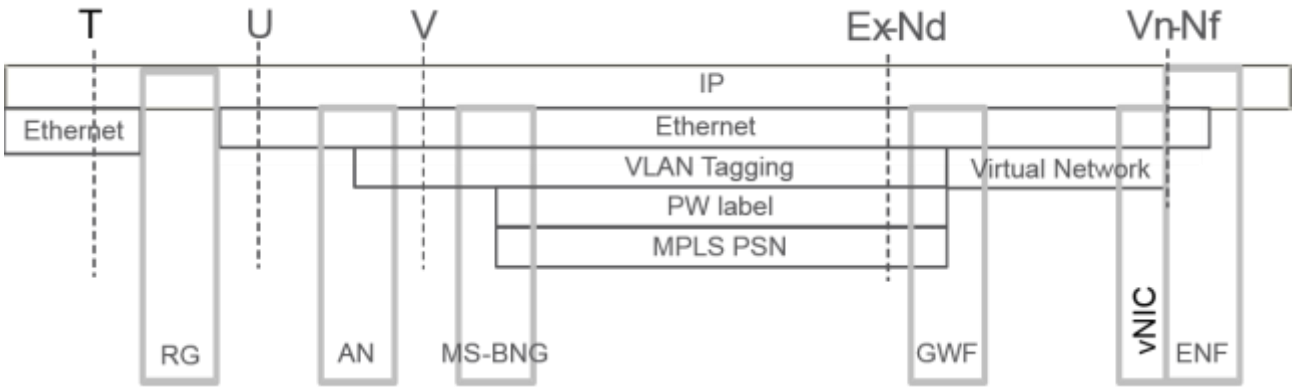

Figure 9 – Protocol stacks for ENF hosted in NFVI – hierarchical MS-BNG scenario

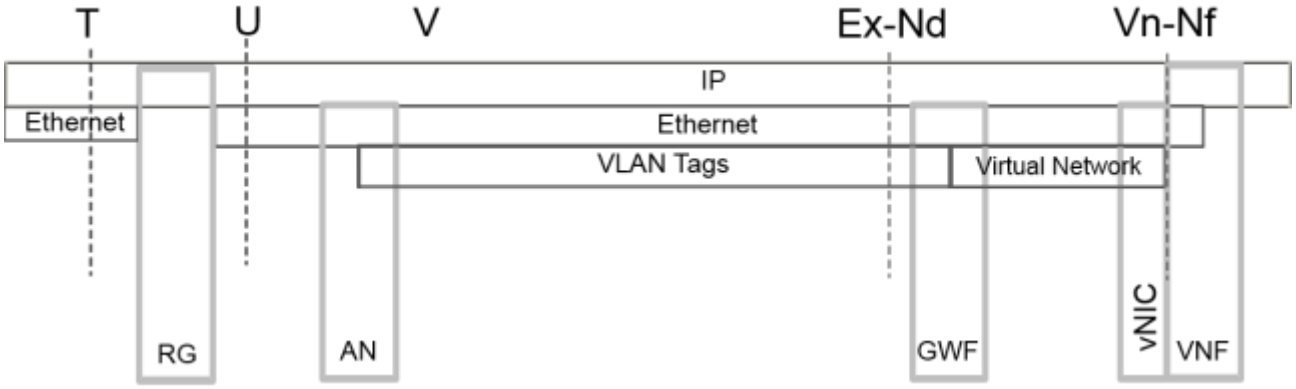

Figure 10 – Protocol stacks for ENF hosted in NFVI – virtual Edge MS-BNG scenario

4.4.3 ENF Hosted on an MS-BNG, with Ethernet or IP Flows presented to a VNF

Figure 11 – ENF Model 2: ENF is hosted on an MS-BNG

Figure 11 shows the scenario where the ENF for the subscriber session is on the MS-BNG and Ethernet or IP Flows are presented to a VNF. There are two models associated with this. The first is an IP presentation of subscriber traffic to subsequent VNFs (see 4.4.3.1) and the second is an Ethernet presentation (see 4.4.3.2).

4.4.3.1 IP Flows

When the model is simply an IP presentation to the NFVI, the Ethernet framing, addressing and tagging is removed by the MS-BNG and the subscriber traffic destined for subsequent VNF processing is tunneled to the appropriate VNF in the NFVI.

The MS-BNG may perform an initial classification of customer traffic into traffic requiring additional processing by functions in the NFVI and traffic which may be directly routed. The initial classification of subscriber traffic uses policy based forwarding as described in section 7.1.1.2 of TR-178.

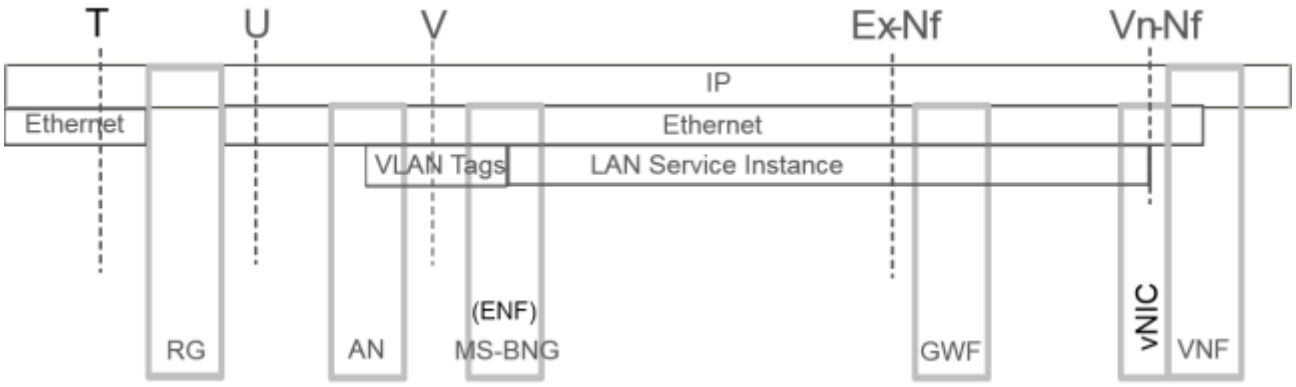

Figure 12 – Protocol stacks for ENF hosted in MS-BNG – IP Presentation

In this model the subscriber session is identified by the source IP address in the subscriber IP packets.

The Ex-Nd interface to the NFVI includes any WAN transport encapsulation added by the MS-BNG.

4.4.3.2 Ethernet Flows

Ethernet frames in the format they are transmitted across the ‘U’ reference point are presented to one or more VNFs in the NFVI across the Vn-Nf interface. This supports use-cases requiring home-LAN extension into NFVI. The ENF in the MS-BNG functions as the AAA enforcement point for session initiation as well as the OAM MEP for the Inter-Carrier ME (the network side logical termination of TR-101 Ethernet access).

The MS-BNG will remove any tagging information and re-map subscriber traffic to a LAN service instance prior to presentation to the Ex-Nd interface to the NFVI. This is done on the basis of the VLAN tag stack at the V-interface.

.

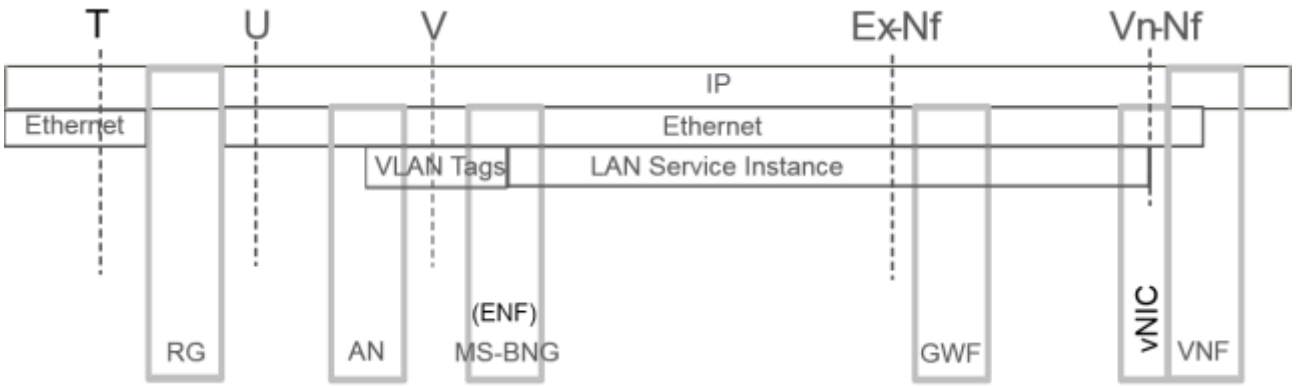

Figure 13 – Protocol stacks for ENF hosted in MS-BNG – Ethernet Presentation

4.4.4 Session Steering

In order to enable service migration or scale ENF performance, multiple ENFs can be provisioned per instance of the V-interface. The access network functions need to achieve a consistent mapping from a subscriber or groups of subscribers to an ENF. This can be a static mapping of a group of subscribers identified by a VLAN stack to an MS-BNG.

A policy interaction on the MS-BNG could instead be used to discover which individual subscriber sessions should be terminated locally on the MS-BNG and which should be mapped to an ENF in the NFVI. In the case of an ENF hosted in NFVI, the MS-BNG or Access Network maps the VLAN stack to an NFVI Gateway Function which then maps the VLAN stack to an NFVI Virtual Network connected to an ENF.

The NFVI Virtual Network can be point-to-point where the final mapping to a VNF is under the control of the NFV Orchestrator. Alternatively, a multipoint Virtual Network can be used which means that the VNFs themselves need to determine which should respond to upstream packets. Where the VNF implements an ENF, an ENF Self-Selection process can use the existing protocol functions of PPPoE, DHCPv4 or DHCPv6 to establish a binding between the subscriber session and the Ethernet MAC address of the ENF.

The ability to “scale-out” or “scale-in” VNF capacity implies the presence of a load-balancing function that influences the mapping of traffic to the first VNF in a service graph. This can be implemented within the NFVI under the control of the VIM. If the MS-BNG incorporates a GWF, mapping of subscriber sessions to Virtual Networks by the MS-BNG could also be part of load-balancing.

4.5 QoS

The H-QoS model utilized in this recommendation adapts the edge and nested BNG H-QoS model from section 4.4.3 of TR-178. The ENF subsumes the H-QoS functionality of a nested BNG. The NFVI-GW / L2PE / L3PE subsumes the H-QoS functionality of the edge BNG but with only a non-subscriber aware classifier for EFP mapping. The ENF performs classification, marking and rate

limiting to provide differentiated QoS capability within the defined envelope of the service rate for the subscriber drop, and communicates that classification to the NFVI-GW / L2PE / L3PE via P-bits and/or DSCP marking.

Such an architecture means that dynamic policy changes are confined to the ENF. More dynamic configurations of H-QoS can be achieved by implementing the GWF integrated with an MS-BNG and leveraging the subscriber and policy management capabilities of TR-178 e.g. enabling the retrieval of scheduling parameters from DHCP and PPPoE line characteristics options and policy server interactions.

5 Use Cases

This section contains a description of the key use cases considered by this Technical Report. These include cases using both Ethernet and IP connectivity to the NFVI. Appendix I shows additional use cases where the virtualization of network functions enabled by TR-345 has applicability to previous Broadband Forum TRs.

5.1 Network Enhanced Residential Gateway

Network Enhanced Residential Gateway (NERG) is an architecture that modifies and enhances the Residential Gateway defined in TR-124, by adding network hosted components. TR-317 defines NERG benefits, use cases, architecture and requirements.

As shown in Figure 14, in the case of NERG, the RG is replaced by the combination of a Bridged Residential Gateway (BRG) and a virtual Gateway (vG). Examples of functions delivered by the vG include IP addressing, Network Address Translation and value added services. The BRG has requirements that are close to an RG defined by TR-124 operating in bridged mode. The vG is a logical component, that may be distributed over multiple nodes, that are physical, virtual or both.

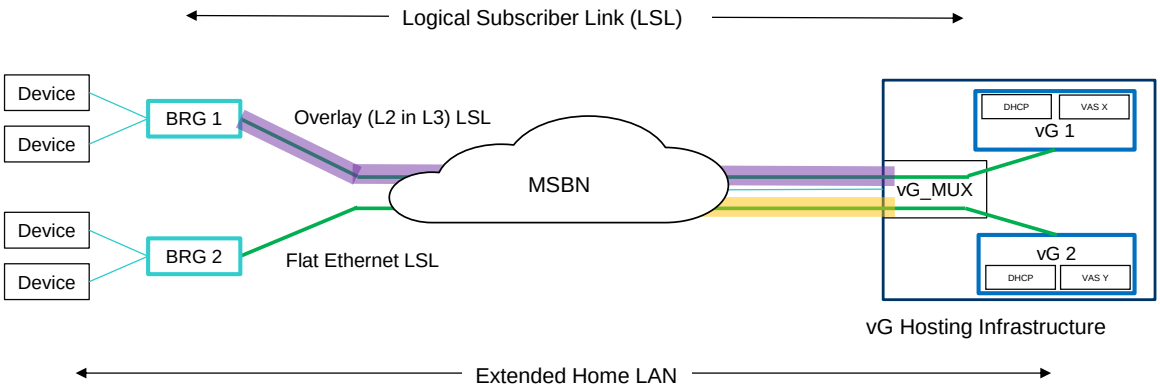

Figure 14 – NERG Logical Components

A BRG and its respective vG are connected by a Logical Subscriber Link (LSL), which carries the subscriber L2 traffic (except the in-home traffic, which is bridged locally by the BRG), with two transport options detailed in this section:

- Flat Ethernet LSL
- Overlay Ethernet LSL

The vG_MUX function selects the proper vG for a given subscriber (note that the LSL terminates at the VG, not at the vG_MUX). Typically, vG_MUX is a function that is hosted either on MS-BNG or on the first VNF in the NFVI.

In the context of TR-345 the vG-MUX is a specialized form of the Classifier Function. The remainder of this section shows examples of the mapping of the LSL and vG hosting infrastructure onto the TR-345 framework.

5.1.1 Flat Ethernet LSL Option

Flat Ethernet means that the L2 traffic is sent natively by the BRG over the WAN interface, without an additional encapsulation (unlike the overlay model, where L2 traffic is encapsulated in an L3 tunnel). In order to maintain segregation between subscribers, this model requires a dedicated VLAN per customer – the access node may add VLAN tags, as shown in Figure 15.

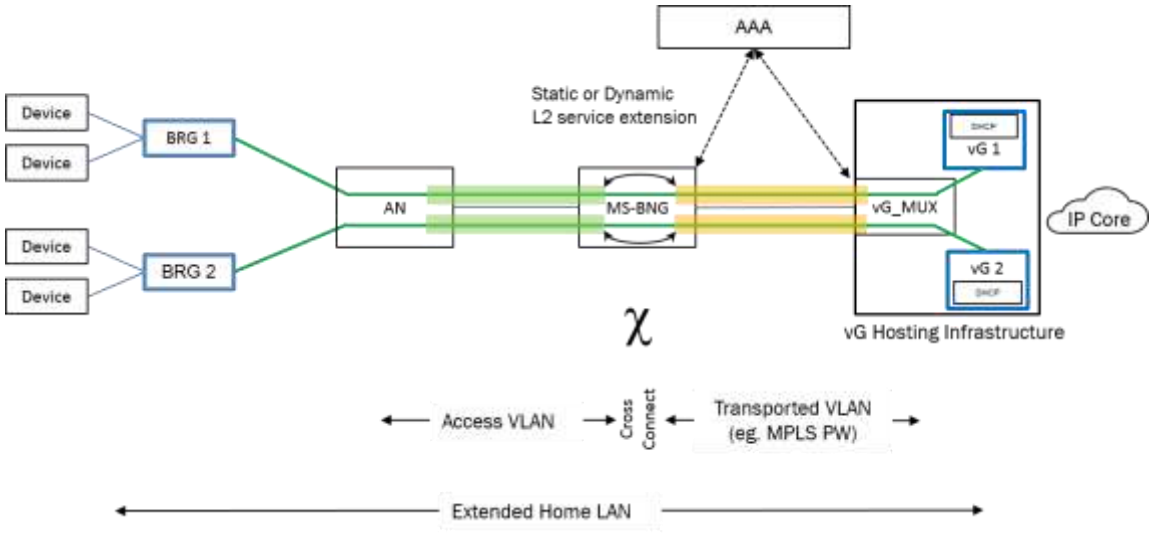

Figure 15 – Flat Ethernet LSL

To extend this subscriber VLAN to the vG, the MS-BNG must either cross-connect VLANs or encapsulate the VLAN for transport, for example with an MPLS PW. The vG_MUX receives the initial Ethernet frame and forwards it to the vG instance for this subscriber, as shown in Figure 16.

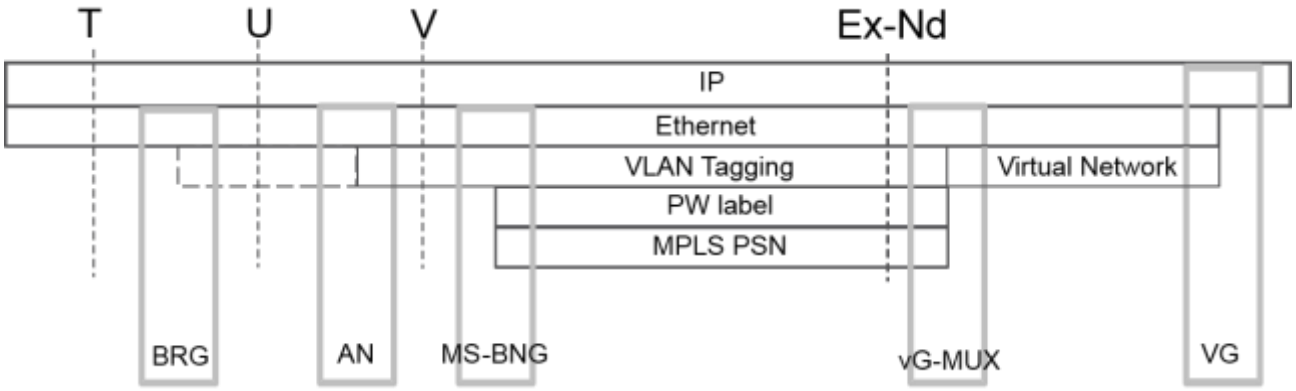

Figure 16 – Packet encapsulations in the case of Flat Ethernet LSL

TR-317 defines static and dynamic ways to establish the LSL. In the dynamic case, AAA can be used to provision the LSL on a per subscriber basis.

The Flat Ethernet LSL model drives specific requirements on the MS-BNG and vG_MUX. In particular, the layer 2 extension on MS-BNG that is specified in TR-178 can be driven by AAA. The vG_MUX is able to terminate MPLS PW and map L2 traffic to a vG instance. Optionally this mapping can be provisioned through a AAA interface.

5.1.2 Overlay LSL Option

5.1.2.1 MS-BNG implements vG Mux and forwards to external vGs

The MS-BNG terminates the GRE-tunnel from the BRG. Each tunnel-session will be steered to different VNF(s) hosting the vGs. The traffic steering is stateful i.e. for the duration of a subscriber session, traffic will always be mapped to the same vG.

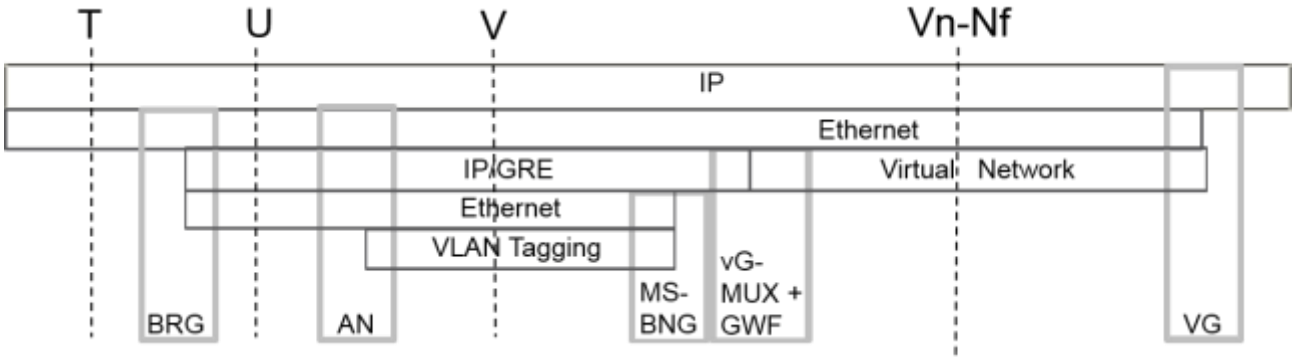

Figure 17 – NERG model with external vGs

The vG_MUX needs to keep subscriber state information for both upstream and downstream traffic. For upstream traffic, the vG_MUX needs to first remove the IP-GRE header and steer the BRG traffic to a dedicated vG over a Virtual Network implemented by the NFVI.

For the downstream traffic, the vG_MUX adds an IP-GRE encapsulation for the GRE-tunnel back to the BRG.

The vG_MUX is also aware of whether or not the vG hosting infrastructure is alive. Upon detecting infrastructurefailure, the vG_MUX must rebalance existing subscriber sessions across the remaining vG hosting infrastructure.

Service providers might consider to firstly adding vG_MUX function to their existing MS-BNGs before virtualizing the vG_MUX function in a later deployment phase.

5.1.2.2 MS-BNG hosts the vG and forwards to external Value Added Services

The MS-BNG will terminate the GRE tunnel from the BRG and host the vG. Traffic will only be steered towards the NFVI for value added services.

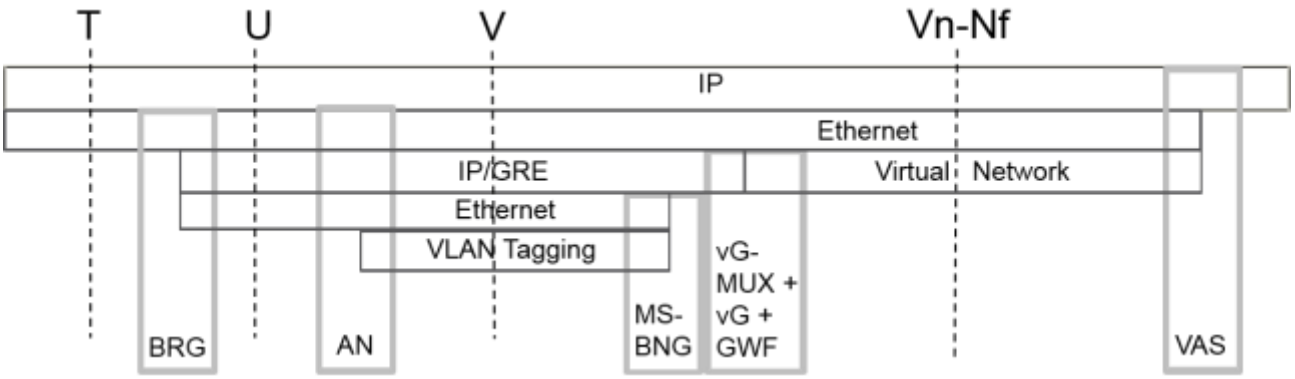

Figure 18 – NERG model with MS-BNG hosting the vG

This use case presents a migration architecture in which service providers reuse their existing network infrastructure to host vGs. In this case the vG will be hosted in the service provider's existing MS-BNGs.

Value added services (VAS), such as firewall and DPI, are VNFs hosted on virtualized platform. For subscribers that have subscribed to these value added services, the vG will steer IP traffic flows from the vG to the appropriate VAS.

5.2 Virtual Business Gateway

The architecture for a Virtual Business Gateway (vBG) is defined by WT-328. The virtual business gateway architecture is similar to the NERG architecture. However, in the Virtual Business Gateway architecture virtual network functions can be hosted on the NFVI at the customer premises in addition to VNFs hosted in the service provider's NFVI. The transport options between the Service Provider NFVI Gateway and the Physical Business Gateway (PBG) remain the same.

5.3 Virtual L2TPv2 Network Gateway

L2TPv2 is used today by network service providers to facilitate broadband services for 3rd parties, in particular:

- Business services, where subscriber sessions are mapped to an enterprise VPN
- Wholesale services, where the LAC resides in the wholesale service provider's network and the LNS in the retail service provider's network.

Virtualizing the LNS allows a full virtual router to be dedicated to each business VPN customer or to each retail service provider (e.g. managed LNS offering). This isolates software management, faults and configurations between one customer and another provided that proper isolation has been put in place and that LNS' are not sharing a common user space. A virtualized LNS also allows granular scale up and down for PPP sessions and for calls per second connected.

The primary business driver is flexibility and agility of deployment by quickly bringing LNS instances up or down, sizing them only to the capacity needed, and configuring them specifically for each customer. That allows a service provider to sell a customized VPN or Retail broadband service

for each customer, sized precisely to required capacity and able to grow quickly as needed. Virtualization can also help address the increased control plane performance requirements due to demanding PPP and L2TP state processing.

The network context is shown here:

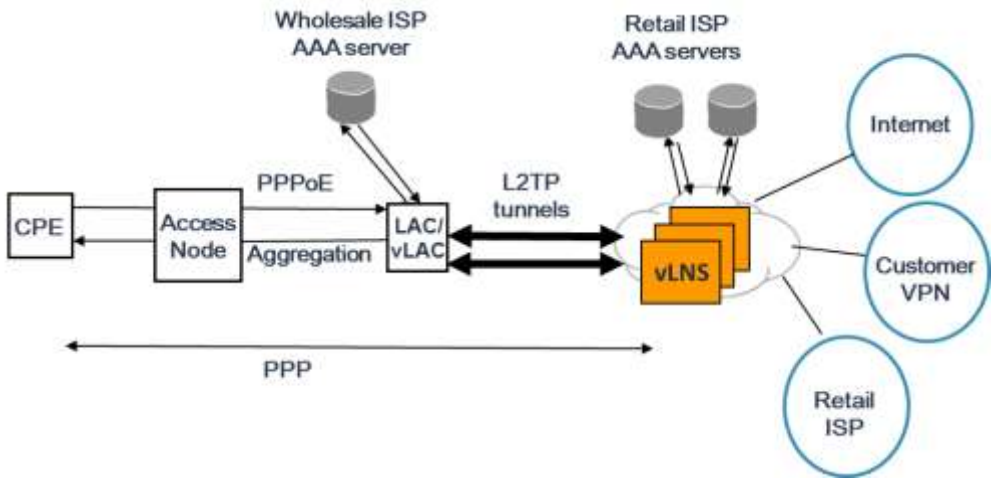

Figure 19 – Virtual L2TP Network Gateway

This is the standard L2TPv2 network environment as it exists today in many wholesale and retail service provider networks, except the LNS is a virtual instance in a datacenter (note that the LAC may also be virtual, but the focus of this use case is the virtualization of the LNS).

The vLNS use case drives the following high level requirements:

- The vLNS IP address needs to be known to the LAC, as it is the destination address of tunneled packets. LAC are configured with LNS information based on either the provisioning system or RADIUS Access-Accept messages. Hence there needs to be some orchestration between the systems that configures the vLNS and the LAC provisioning system or the RADIUS server attached to the LAC.
- The NFVI Gateway Function needs to accommodate L3 connectivity between the LAC and the vLNS (routable IP address).
- Load-balancing and resiliency will continue to be based on LNS selection provisioned at the LAC level. To leverage the scaling agility provided by NFV, this selection logic needs to be aware of the vLNS instance scaling properties.
- The vLNS needs to be reachable from the retailer RADIUS server. The vLNS IP address can be be provisioned in the RADIUS server, but this is optional.
- The vLNS instances needs to be able to put core facing traffic in an L3 MPLS routing instance for business services. This might require L3 MPLS support in the gateway routers.

6 Node and Virtual Function Requirements

This section provides requirements for the MS-BNG, NFVI Gateway and Virtual Functions identified in section 4.

6.1 MS-BNG Requirements

6.1.1 Session Control

The MS-BNG needs to determine how to forward the traffic for each subscriber session. This traffic could either be forwarded natively by the MS-BNG or steered towards the NFVI. TR-178 contains requirements for subscriber management that allows the rules for each subscriber session to be determined through policy lookups triggered on first sign of life. These capabilities are reused by this Technical Report.

[R-1] The MS-BNG MUST support subscriber management functions in TR-178 section 7 (Authentication, Accounting, policy-enforcement, subscriber-management, and address assignment). **M, C**

[R-2] The MS-BNG MUST be able to update traffic steering policies for existing subscriber sessions based on policy server requests.

6.1.2 Traffic Steering

Traffic steering is performed by the Classifier Function in the MS-BNG.

[R-3] The MS-BNG MUST support traffic steering on a per 5-tuple basis. **M, C**

[R-4] The MS-BNG SHOULD support traffic steering on a per application basis. This relies on mechanisms such as Layer 7 classifiers e.g. HTTP URL. **M, C**

[R-5] The MS-BNG MUST be able to identify IP flows and steer them to bypass the NFVI. **M, C**

[R-6] The MS-BNG MUST NOT depend on C-VLAN-ID uniqueness for distinguishing subscribers

6.1.3 QoS

In the scenario where the ENF is implemented within the MS-BNG, the MS-BNG hierarchical scheduling function performs hierarchical scheduling as per requirements documented in TR-178.

6.1.4 Embedded NFVI Gateway Function

An MS-BNG with an embedded NFVI Gateway Function can interwork the protocols used on access facing interfaces with those used by the NFVI.

[R-7] An MS-BNG supporting an Embedded NFVI Gateway MUST support the NFVI Gateway Function requirements in section 6.3.1.

6.2 Standalone NFVI Gateway Requirements

Virtualized network functions are reachable via an NFVI Gateway by the use of either bridging of Ethernet frames, or forwarding of IPv4 or IPv6 packets depending on the network functions deployed.

When a standalone NFVI-GW is deployed it is required to support interconnection with existing BBF specified networks.

This section describes the scenarios supported by BBF TRs and their associated requirements. More detail on the architectures driving these scenarios is given in Appendix I.

6.2.1 NFVI-GW directly connected to a TR-101 network supporting a virtualized MS-BNG

In this model there is no physical MS-BNG and the NFVI-GW is directly connected to a TR-101 based access network. For a given subscriber, the NFVI-GW bridges traffic to first ENF instance in the virtualized MS-BNG.

[R-8] The NFVI-GW MUST support requirements R1 through R28 (NLIM requirements) in TR-178.
[R-9] The NFVI-GW MUST support requirements R300 to R304 (Traffic Management) in TR-178. **M**
[R-10] The NFVI-GW MUST support all requirements in section 7.1.3 (OAM) in TR-178. **M**
[R-11] The NFVI-GW MUST be able to perform EFP classification in order to map into H-QOS queues. **M**
[R-12] The NFVI-GW MUST support all requirements in section 7.1.7 of TR-178. **M**
[R-13] The NFVI-GW MUST support the NFVI Gateway Function requirements in section 6.3.1.

6.2.2 L2PE supporting NFVI hosting virtualized MS-BNGs or NERG in a flat LSL model

In this model existing TR-178 mechanisms are used to tunnel TR-101 based Ethernet subscriber traffic from an edge deployed MS-BNG to the GWF where it is de-multiplexed to the VNF instance(s) implementing the virtualized MS-BNG or NERG. Therefore the GWF is required to terminate the EPL/EVPL service, and decapsulate and bridge the subscriber traffic to the appropriate ENF.

[R-14] The L2PE MUST support requirements R300 to R304 (Traffic Management) in TR-178. **M**
[R-15] The L2PE MUST support all requirements in section 7.1.3 (OAM) in TR-178. **M**
[R-16] The L2PE for Ethernet Private Line (EPL) MUST support all requirements in Section 11.2.1/TR-224 [8]. **M**

[R-17] The L2PE for Ethernet Virtual Private Line (EVPL) MUST support all requirements in Section 11.3.1/TR-224. **M**
[R-18] The L2PE for EPL or EVPL MUST support the CoS requirements in sections 9.1 and 9.2 of TR-224. **M**
[R-19] An L2PE SHOULD support the EVPN functionality specified in TR-350 section 11. **M**
[R-20] The L2PE MUST support the NFVI Gateway Function requirements in section 6.3.1.
[R-21] The L2PE MUST support the MPLS Termination requirements in section 6.3.2.

6.2.3 L3PE supporting NFVI hosting BSG functions, L2TP functions, a 3GPP TWAG, or IPv4-IPv6 migration functions

Within the Broadband Forum Technical Reports there are a number of defined functions that operate on traffic tunneled to them from either a MS-BNG or an RG. This requires a GWF to be able to steer or route traffic to a VNF on the basis of IP address. A non-exhaustive list of IP encapsulation mechanisms is:

- GRE tunnel from an RG to a TWAG VNF
- GRE or other tunnel from an RG to a NERG vG_MUX
- L2TP tunnel from a LAC to a L2TS or LNS VNF
- IPv4inIPv6 DSLite tunnel to a DSLite AFTR
- IPv6inIPv4 6RD tunnel to a 6RD BR VNF

A GWF in this deployment scenario supports Layer 3 VPN services to map external traffic into the NFVI.

[R-22] An L3PE providing external connectivity for VNFs that terminate MPLS encapsulated L3 tunnels MUST support the MPLS PE functionality specified in TR-221 section 8.4. **M**
[R-23] The L3PE MUST support the NFVI Gateway Function requirements in section 6.3.1.
[R-24] The L3PE MUST support the MPLS Termination requirements in section 6.3.2.

6.3 Virtual Function Requirements

6.3.1 NFVI Gateway Function Requirements

6.3.1.1 Protocol Interworking

The NFVI Gateway Function (GWF) is common to a number of platform instantiations. The GWF maps traffic from the external network to NFVI L2 and/or L3 networks.

When the ENF is in the NFVI, the GWF maps a VLAN stack received from a TR-101/TR-178 access network to a broadcast domain in the NFVI that connects the GWF to the ENF system.

[R-25] The GWF MUST be able to map the S-VLAN-ID or S-VLAN-ID/C-VLAN-ID tuple in a received frame to an NFVI broadcast domain that interconnects the subscriber facing interfaces of an ENF System to the GWF. **M**

[R-26] The GWF MUST be able to remove and re-apply the S-Tag or S-Tag/C-Tag tuple for traffic received from/relayed to the subscriber. **M**

[R-27] A GWF supporting TR-317 overlay LSLs and supporting vG-Mux functionality MUST implement the GRE tunnel procedures in section 5.1.2. **M**

When the ENF is in an MS-BNG, the GWF maps an IP Flow to a VNF in the NFVI.

[R-28] The GWF MUST be able to map an IP Flow to an NFVI L2 or L3 Virtual Network. **M**

6.3.1.2 Load Balancing

The GWF may embody a load balancing function configured by the VIM, otherwise a load balancer within the NFVI can be used. VNF bandwidth (compute-based) can be much lower than WAN interface bandwidth. A load balancer in the GWF can be used to distribute traffic for a logical function between multiple VNF instances. Stateful services supporting functions such as DPI or parental control service functions can require load balancing to be deterministic to minimize requirements for state sharing between VNFs.

[R-29] The GWF SHOULD support load-balancing across a set of VNF's in order to support both scaling and redundancy. **M, C**

[R-30] The GWF SHOULD support deterministic load-balancing, so that specific flows are steered to the same VNF, in both directions (for example, CGNAT). **M, C**

[R-31] The GWF load balancing function MUST maintain a binding between a subscriber session and a VNF.

[R-32] After a failure or removal of a VNF, the load balancing algorithm MUST redistribute the subscriber sessions mapped to the failed VNF without changing the mapping of subscriber sessions mapped to surviving VNFs.

[R-33] The addition of a VNF to a load balanced set MUST not affect the mapping of existing subscriber sessions to VNFs.

[R-34] The GWF MUST only steer traffic to NFVI platforms known to be live.

6.3.2 MPLS Termination

An L2PE / L3PE that supports MPLS termination as per section 6.2.2 or 6.2.3 will need to support a number of MPLS functions to connect to an MPLS network:

[R-35] A PE MUST support all the signaling & routing requirements specified in TR-224 section 7. **M**

[R-36] A PE MUST support all the MPLS OAM requirements specified in TR-224 sections 8.2.1 to 8.2.4. **M**

[R-37] A PE MUST support all the resiliency requirements of TR-224 section 10.3. **M**

6.3.3 NFVI Network Domain Requirements

A number of implementations of NFVI Virtual Networks depend a management system being able to control and pre-provision end system MAC addresses in the forwarding tables of virtual switches (e.g. D-MAC to tunnel bindings). The implication is that NFVI does not offer the traditional "plug

and play" operation normally associated with Ethernet LANs. Therefore some form of interworking is required between a "plug and play" LAN or RG and the NFVI network supporting the ENF system.

[R-38] The NFVI Network Domain MUST support L2 connectivity to end systems where the MAC addresses are not administered by the VIM and where such systems utilize traditional bridging.

[R-39] The NFVI Network Domain MUST support L2 broadcast.

6.3.4 ENF Requirements

The requirements in the section apply only to a virtual ENF hosted in NFVI. Requirements for the ENF in an MS-BNG are specified by existing Broadband Forum Technical Reports (e.g. TR-178 and TR-146).

6.3.4.1 ENF Self Selection

ENF Self Selection is an option for the steering of subscriber sessions to an ENF once they have been mapped to a Virtual Network by a GWF.

When the ENF is hosted in the NFVI, the mapping between a set of subscribers and an ENF can be pre-provisioned or dynamically chosen by a virtual MS-BNG. In the latter case, in order to simplify the provisioning of the GWF, the suggested technique for the distribution of load across a set of virtualized ENFs depends on existing protocol features that exist in both the PPPoE and IPoE protocol suites.

The set of ENF instances that serves one or more S-VLAN delineated broadcast domains is an ENF System.

The initial broadcast message from the subscriber that initiates a subscriber session (PPPoE or IPoE) is responded to by the ENF system such that it will ultimately resolve to the MAC address or addresses of the ENF instances that can host the subscriber session.

When the subscriber end system receives offers (PADI or DHCP Offer) that correspond to more than one ENF instance, the subscriber end system chooses which offer to accept and proceeds with session initiation.

The management of the ENF system may direct only a subset of ENF instances to be identified in the initial offers in order to control the distribution of subscriber load across the set of ENF instances.

The overall implementation can be scaled by the provisioning of multiple ENF systems such that the number of subscribers served by any single broadcast domain is bounded. This reduces the broadcast of unknown MAC addresses and limits the number of MAC addresses that need to be learned within a single broadcast domain.

[R-40] The ENF system SHOULD appear as a single AAA client.

[R-41] The PPPoE ENF system MUST provide at least one PADO offer in response to a PADI request.

[R-42] The MAC address used as the MAC Source Address in the PADO response MUST correspond to the ENF instance that will host the subscriber's IP session.

[R-43] The ENF system SHOULD ensure that all subscriber IP sessions associated with a common Circuit ID and Remote ID are hosted on a common ENF instance.

Note: This addresses having a common ENF for multiple bridge subscriber end systems. This has no impact on TR-187 [6] single PPP, dual stack sessions. Both IPv4 and IPv6 would be served by a common ENF instance.

[R-44] The IPoE ENF system MUST provide at least one DHCP Offer in response to a DHCP Discover request.

[R-45] The IPoE ENF system MUST provide at least one DHCPv6 Advertisement in response to a DHCPv6 Solicit request.

[R-46] The GIADDR field in the DHCP Offer response MUST uniquely identify the ENF instance in the ENF system that will host the subscriber's IP session.

[R-47] The ENF System MUST ensure that only the ENF instance that will host the subscriber's IPv6 session offers an IPv6 RA in response to an RS.

Note: It is advised that operational practice should ensure that the NFVI operations do not partition the ENF system outside a single NFVI site. Further the ENF system should hide artifacts of VM migration from the routing system.

[R-48] The ENF system MUST be able to reply to ARP messages for every GIADDR address offered to the served set of subscribers

There are three ways in which a line ID can be expressed by transactions originating from the customer premises. This is via:

- Option 82 insertion by an AN located DHCPv4 relay agent
- Option 18 insertion by an LDRA
- RFC-6788 line ID insertion in IPv6 RS messages [19]

[R-49] The ENF system MUST ensure that all IPv4 and IPv6 sessions that are associated with subscriber (identified by a common line ID,VLAN tag stack or MAC address) are served by a common ENF instance.

[R-50] The IPoE ENF system that supports IPv6 MUST provide at least one RA message in response to a RS

[R-51] RA messages issued by the ENF system MUST be sent as Ethernet unicast packets as per RFC 6085 [18] / TR-177 R44 [4].

6.3.4.2 Traffic Management and QoS

In the scenario where the ENF for the IP session is deployed in the NFVI, the ENF includes a traffic management function (TMF) that both rate limits the aggregate of downstream traffic transiting the EFP and provides differential queuing. The rate limit value used is a configured rate and may be dynamically modified via IGMP joins/leaves to adjust for any AN inserted multicast streams as described in TR-101.

[R-52] An NFVI hosted ENF MUST support per EFP rate limiting **M, C**

[R-53] An NFVI hosted ENF MUST support traffic classification. **M, C**

[R-54] The ENF MUST be able to assign a minimum bandwidth guarantee to a queue. **M, C**

Note: the minimum can be zero (no guarantee).

[R-55] The ENF MUST be able to map traffic classes to queues. These queues MUST be able to be configured with a given forwarding behavior such as Expedited Forwarding, Assured Forwarding, etc. **M, C**

[R-56] Per Class Queuing of traffic MUST be supported within the ENF's EFP rate limited instance. **M, C**

[R-57] The ENF MUST be able to be configured to terminate IGMP / MLD traffic in order to track AN multicast insertion. **M, C**

Note: This can be used to dynamically update the user-facing QoS shapers per the requirements in TR-101 section 6.3.2.3.1.

Appendix I. Comprehensive Infrastructure Model

This appendix expands on the use cases described in section 5 by cataloguing a broader set of applications for tunneling within Broadband Forum architectures. If a function is already defined as a tunnel end-point, then by extension it can be separated from the MS-BNG and a candidate for migration to Network Function Virtualization.

Figure 20 illustrates the set of network functions that have been defined in BBF Technical Reports. For a number of these functions tunneling is used to either access them or as an inherent part of the function offered. Such tunnels may originate at an RG or an edge deployed MS-BNG. This occurs for numerous reasons:

1. Wish to operationally decouple the provisioning of tunneled connectivity from provisioning of subscriber session termination on the MS-BNG.
2. Wish to extend TR-101 backhaul beyond the Edge MS-BNG.
3. Support IPv6 migration strategies

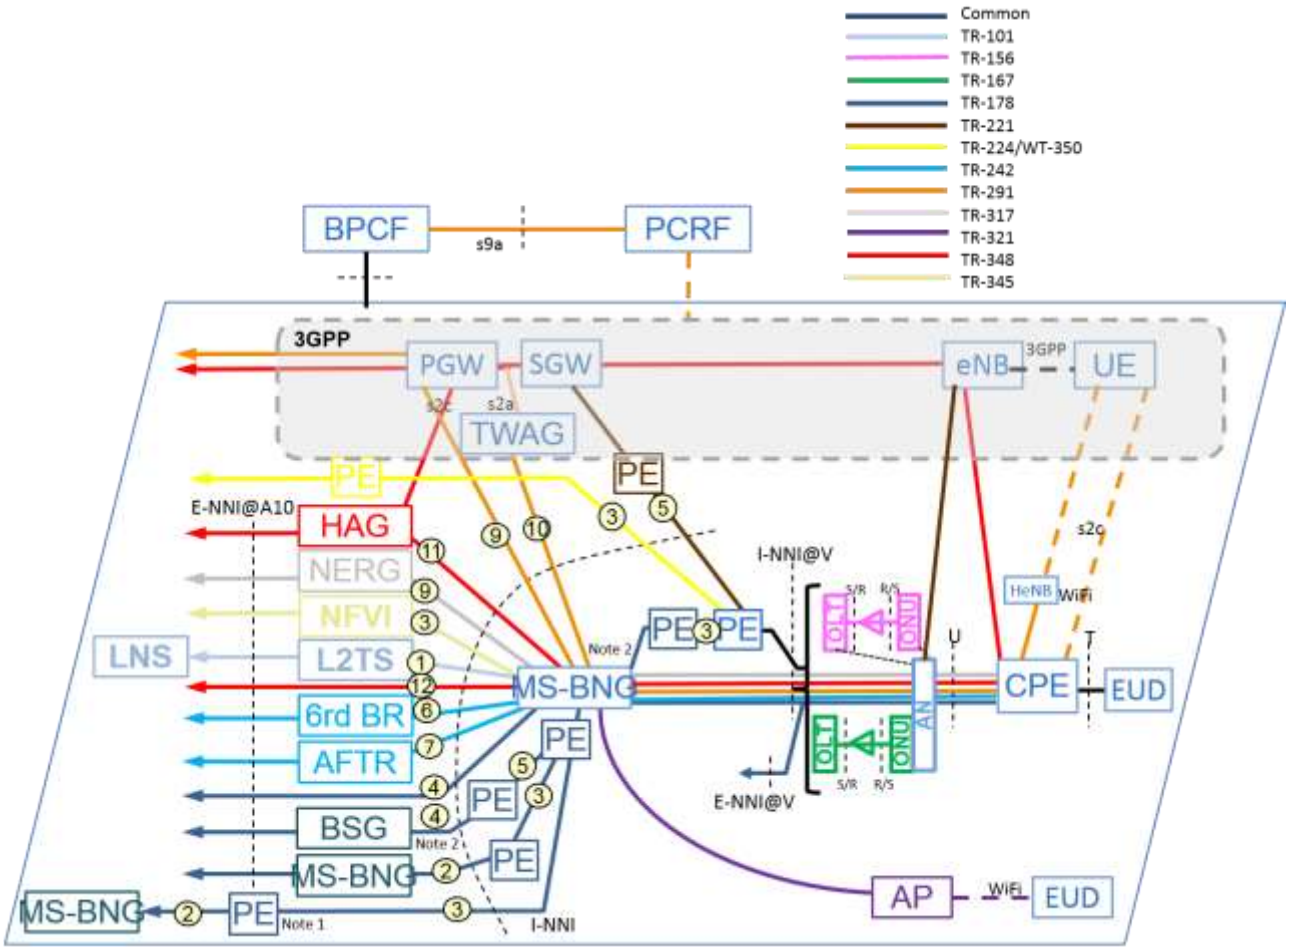

AFTR – Address Family Transition Router
AP – Access Point
BPCF – Broadband Policy Control Function
BR – Border Router
BSG – Broadband Service Gateway
CPE – Customer Premises Equipment
eNB – Evolved Node B
EUD – End User Device
HAG – Hybrid Access Gateway
HeNB – Home eNB
L2TS – Layer 2 Tunnel Switch
LAC – L2TP Access Concentrator
LNS – L2TP Network Server
MS-BNG – Multiservice Broadband Network Gateway
NERG – Network Enhanced Residential Gateway
NFVI – Network Functions Virtualisation Infrastructure
OLT – Optical Line Termination
ONU – Optical Network Unit
PE – Provider Edge
PGW – P-Gateway
SGW – S-Gateway
TWAG – Trusted WLAN Access GW
UE – User Equipment

Note 1: PE can be on either side of the A10
Note 2: In TR-178 PE←→MS-BNG interfaces may be NULL

Figure 20 – Comprehensive Infrastructure Model

The following table illustrates the stacks in use indexed by the particular function and cross referenced to the relevant TR. The table indicates the protocols in use at the U and I-NNI interfaces, and to assist comprehension identifies if a tunnel is used, whether it originates at the RG or the MS-BNG.

Table 1 – Distinct Protocol Models

#	TR	Tunnel Origin	U Reference Point	I-NNI Reference Point
1	TR-101	MS-BNG	PPPoE[3]	PPP/L2TP/IP/Foo
2	TR-178	N/A	PPPoE, IPoE	Ethernet
3	TR-178 TR-224 TR-345 TR-350	MS-BNG	Foo/Ethernet	IpoEor PPPoE/PW/MPLS/Foo Foo/Ethernet/PW/MPLS/Foo IpoEor PPPoE/PW/MPLS/Foo Foo/Ethernet/MPLS/Foo
4	TR-178	N/A	PPPoE, IpoE	IP/Foo
5	TR-178 TR-221	N/A	PPPoE, IpoE IpoE	IP/MPLS/foo
7	TR-242	RG	Ipv6/Ipv4/PPPoE or Ipv4oE	Ipv6/Ipv4/foo
8	TR-177/187 TR-242	RG	Ipv4/Ipv6/PPPoE or Ipv6oE	Ipv4/Ipv6/foo
9	TR-291 TR-317 TR-321	RG	Ethernet/GRE/IP-PPPoE or IpoE	Ethernet/GRE/IP/foo
10	TR-291	RG	s2c GTP/IPSEC/PPPoE, IpoE	s2c GTP/IPSEC/foo
11	TR-348[2]	RG	TCP/MP-TCP/IP/PPPoE or IpoE	TCP/MP-TCP/IP/foo
12	TR-348[2]	RG	UDP/IP	UDP/IP/foo

Note 1 - Foo = unspecified lower layer protocol(s)

Note 2 - TR-348 may ultimately describe multiple encapsulations

Note 3 - PPPoE identified in isolation as IPoE does not map to L2TP at the I-NNI

End of Broadband Forum Technical Report TR-345