TECHNICAL REPORT

TR-317
Network Enhanced Residential Gateway

Issue: 1
Issue Date: July 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	18 July 2016	5 August 2016	David Minodier, Orange Gregory Dalle, Juniper Networks	Original

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

Editors	David Minodier Gregory Dalle	Orange Juniper Networks	david.minodier@orange.com gdalle@juniper.net
Architecture & 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

Executive Summary

TR-317 specifies the Network Enhanced Residential Gateway (NERG) architecture. NERG consists in shifting some of the functionality of a residential gateway (RG), as defined in TR-124 [6], to the operator's network.

With NERG, the functions provided traditionally by the RG are now distributed between ant on-site device called a BRG (Bridged Residential Gateway) and a network based component, called a vG (virtual Gateway). The vG is part of a flexible hosting environment that can benefit both from network equipment and recent network virtualization technology.

This Technical Report describes the motivation to deploy NERG, based on the use cases that it enables. In particular, NERG facilitates per subscriber per device services, LAN extension and simplified end user operations. Service providers looking for additional service offerings, such as personalized services, Smart Home and enhanced multimedia services, may find significant benefits in deploying NERG.

Following a high level architecture description and some examples of deployment models, TR-317 specifies the following set of technical requirements:

- End-to-end network requirements and support for existing, as well as new, broadband services
- BRG requirements, referencing TR-124 wherever possible
- vG requirements

The target audience for this document is:

- Network service providers who want to evaluate the impact of NERG on their services and architectures,
- Suppliers who want to build interoperable BRGs and vGs,
- System integrators responsible for the integration of NERG in NSP’s information systems

1 Purpose and Scope

1.1 Purpose

This Technical Report specifies the Network Enhanced Residential Gateway (NERG) architecture. This architecture consists in shifting some of the functionalities of a Residential Gateway to the operator's network, to enable network-based features. The aim is to facilitate the deployment, maintenance and evolution of both existing and new capabilities without adding complexity to the Residential Gateway and/or the home network.

1.2 Scope

The scope of this Technical Report is:

- Use cases and business drivers for moving specific Residential Gateway functions to the network
- High level network architectures to support NERG
- The nodal requirements for the customer located Bridged Residential Gateway (BRG) and the network component called the virtual Gateway (vG), in order to support interoperable NERG implementations
- The connectivity between the BRG and the vG, including related AAA requirements
- The role of the NERG in IPv4 address management and migration to IPv6
- Requirements for supporting both new and existing residential services (Internet access, IPTV, VoIP, VoD, WiFi) with these architectures.
- QoS
- Ensuring continuing home network availability, in the case of failure between the BRG and the vG
- Management of NERG functions
- Security
- Privacy

In this document, the vG is defined as a logical component. Although its main functional subcomponents are described and specified, their nodal distribution within the vG and their interconnections are not covered. In particular, the following aspects are out of scope:

- Connecting virtual network functions to a TR-178 broadband network – please refer to WT-345 [18] and WT-359 [19]
- Service chaining and SDN.

The following areas are also out of scope.

- Over-the-top NERG, where the NERG components would be operated by an administrative entity other than the NSP operating the broadband network
- PPP support (either in the BRG or vG)
- Per device services and policy management (e.g. per device QoS or steering)
- Support of more than one BRG by the same vG - this document assumes a 1:1 relation between BRG and vG.
- NERG orchestration

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 [5].

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.

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-069 Amendment 5	*CPE WAN Management Protocol*	BBF	2013
[2]	TR-101 Issue 2	*Migration to Ethernet-Based Broadband Aggregation*	BBF	2011
[3]	802.1D	*MAC bridges*	IEEE	1993
[4]	802.1Q	*Virtual LANs*	IEEE	2006

[5]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[6]	TR-124 Issue 4	*Functional Requirements for Broadband Residential Gateway Devices*	BBF	2014
[7]	ETSI GS NFV 0002 V1.1.1	*Network Functions Virtualization (NFV); Architectural Framework*	ETSI	2013
[8]	RFC 2784	*Generic Routing Encapsulation (GRE)*	IETF	2000
[9]	RFC 2868	*Zorn et.al., "RADIUS Attributes for Tunnel Protocol Support"*	IETF	2000
[10]	RFC 3550	*RTP: A Transport Protocol for Real-Time Applications*	IETF	2003
[11]	RFC 3580	*IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines*	IETF	2003
[12]	RFC 4861	*Neighbor Discovery for IP version 6*	IETF	2007
[13]	RFC 6731	*Improved Recursive DNS Server Selection for Multi-Interfaced Nodes*	IETF	2012
[14]	RFC 6887	*Port Control Protocol (PCP)*	IETF	2013
[15]	TR-144	*Broadband Multi-Service Architecture & Framework Requirements*	BBF	2007
[16]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[17]	TR-242i2	*IPv6 Transition Mechanisms for Broadband Networks*	BBF	2015
[18]	WT-345	*Migrating to NFV in the context of WT-178*	BBF	2016
[19]	WT-359	*Framework for Virtualization*	BBF	2016
[20]	TR-146	*Subscriber Sessions*	BBF	2013

2.3 Definitions

The following terminology is used throughout this Technical Report.

Subscriber	The customer from a contractual perspective.
User/end-user	A person who uses a service.
Administrator	User who has some privileges for the configuration, management and diagnosis of the services he purchased from the NSP.
NERG	Network Enhanced Residential Gateway. Distributed RG where the RG functions are distributed between the BRG at the customer premises and the vG in the network.

NERG Session Ethernet Session which represents all traffic transported between the BRG and the vG over the LSL, these three elements being associated with a given subscriber.

Actor A participant in one or more of the management, operation and use of a NERG.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AAA	Authentication, Authorization and Accounting Server
ACS	Auto-Configuration Server. This is the component in the broadband network responsible for auto-configuration and troubleshooting of the RG.
AF	Assured Forwarding
ALG	Application Layer Gateway
ATA	Analog Telephone Adapter
BE	Best Effort
BRG	Bridged Residential Gateway
BSS	Business Support System
B-DHCP	Backup DHCP server
B-DNS	Backup DNS server
CGN	Carrier Grade NAT
COTS	Commercial Off The Shelf
CPE	Customer Premises Equipment.
DDoS	Distributed Denial of Service (attack)
DHCP	Dynamic Host Configuration Protocol
DHCPv6	Dynamic Host Configuration Protocol for IPv6
DHCPv6-PD	Dynamic Host Configuration Protocol for IPv6 with Prefix Delegation
DLNA	Digital Living Network Alliance
DNS	Domain Name Service
DPI	Deep Packet Inspection
E2E	End-to-End
EF	Expedited Forwarding
EMS	Element Management System
FTP	File Transfer Protocol
FTTH	Fiber To The Home
FQDN	Fully Qualified Domain Name

GRE	Generic Routing Encapsulation
GUI	Graphical User Interface
HPNA	Home Phone line Networking Alliance
IA_NA	Identity Association for Non-temporary Addresses
IGD	Internet Gateway Device
IP	Internet Protocol
IPTV	Internet Protocol Television
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
IT	Information Technology
IWF	Interworking Function
LAN	Local Area Network
L2	Layer 2
L2TPv3	Layer 2 Tunneling Protocol version 3
LSL	Logical Subscriber Link
M2M	Machine-To-Machine
MoCA	Multimedia over Coax
MEP	Maintenance association End Point
MS-BNG	Multi-Service Broadband Network Gateway
MSBN	Multi-Service Broadband Network
NAS	Network Attached Storage
NAT	Network Address Translation
NERG	Network Enhanced Residential Gateway
NF	Network Function
NFV	Network Function Virtualization
NFVI	Network Function Virtualization Infrastructure
NMS	Network Management System
NSP	Network Service Provider
NTP	Network Time Protocol
OAM	Operations Administration and Maintenance
OSS	Operational Support System
OTT	Over The Top
PCP	Port Control Protocol
PMP	Port Mapping Protocol
PNF	Physical Network Function

PoP	Point of Presence
PW	Pseudo Wire
QoS	Quality of Service
RADIUS	Remote Authentication Dial in User System
RG	Residential Gateway
RTCP	Real-time Transport Control Protocol
RTP	Real-time Transport Protocol
RTSP	Real Time Streaming Protocol
SDN	Software Defined Network
SDP	Session Description Protocol
SLAAC	Stateless Address Auto-configuration
STB	Set Top Box
S-VLAN	Service VLAN
TMF	Traffic Management Function
ToR	Top of Rack
TR	Technical Report
TTM	Time To Market
UI	User Interface
UPnP	Universal Plug and Play
vG	Virtual Gateway
vG_MUX	Virtual Gateway Multiplexer
vNAS	virtual Network Attached Storage
VAS	Value Added Service
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
VNI	Virtual Network Identifier
VoIP	Voice Over IP
VPWS	Virtual Private Wire Service
VxLAN	Virtual eXtensible Local Area Network
WAN	Wide Area Network
WLAN	Wireless Local Area Network.
WT	Working Text

3 Technical Report Impact

3.1 Energy Efficiency

The NERG architecture relocates some RG functionalities in the network. This will result in some increased consumption of energy in the service provider network. However, it is believed that the overall efficiency of the system (premises plus network located functions) could be increased if functionalities provided by CPEs other than the RG are also moved to the service provider network.

3.2 IPv6

The NERG architecture facilitates the migration to IPv6 as described in section 5.5 for example providing a NAT function as a component in the vG and allowing the reconfiguration of non-IPv6 capable RGs as a L2 BRG, so the RG itself does not need to be made IPv6 aware.

3.3 Security

TR-317 does have an impact on security. Various aspects of a single dedicated system with a routed RG on the customer premises will now be implemented in a number of network hosted functions. This does increase the number of potential avenues for malicious attack, as most of the systems implementing a vG will utilize shared resources, the vulnerability to DOS attacks in particular may increase. The provider of the network aspects of the NERG will need to address these issues.

3.4 Privacy

The NERG architecture does expose the home devices to the operator's network – they are no longer hidden behind an on premise NAT function – which may lead to some privacy concerns. This section compares legacy RGs to the NERG architecture with regards to privacy.

3.4.1 Privacy with a routed RG

A managed RG is controlled by the NSP. In this model, the operator can already access detailed information about the home LAN devices. The NSP has root access to the RG and can therefore activate some level of DPI in the RG, collect statistics, or even mirror some LAN traffic.

A non-managed RG is not controlled by the operator but does not typically support the triple or quadruple-play service bundles offered by the NSP.

Subscribers concerned with privacy may add their own non-managed router behind the service provider's RG in order to hide their LAN devices from the operator. However, this would not prevent the operator from monitoring the outgoing traffic to the Internet and HTTP headers contain metadata providing information on devices (browser type, version, screen resolution, etc.).

3.4.2 Privacy with a BRG and NERG

3.4.2.1 Device visibility

LAN devices connected to a BRG have the same exposure to the operator as those connected to a managed legacy RG. However with NERG, the NAT and firewall functions are now located in the network. Therefore the vG should be hosted in secure premises such as a PoP or Datacenter.

3.4.2.2 Device accessibility

Just as in the legacy architecture, access to LAN devices from the Internet and from other subscribers is prevented by the vG NAT and firewall functionality in the case of IPv4, and the vG firewall alone for IPv6.

NERG subscribers concerned with privacy can insert their own router behind the BRG to hide their devices, similar to the RG case. However this would prevent the subscriber accessing most of the NERG-specific services.

3.4.2.3 Local traffic

Traffic between two LAN devices located in the home remains local since the BRG acts as an Ethernet switch capable of doing MAC Learning. Only unicast traffic to be sent to the Internet, broadcast and multicast messages are forwarded to the vG (DHCP, DLNA, ARP, etc.). Local unicast traffic is not sent to the network (file transfer between two computers, video streaming from a local NAS to a TV set, printing, etc.).

3.4.2.4 Conclusion

The NERG architecture is similar to the managed RG, legacy architecture with respect to privacy. However as routing, NAT and firewall functions are moved to the network, it is important that the vG be located in secure premises. Wary subscribers can add their own router behind their bridged or managed routed residential gateway to hide all or a subset of their devices.

4 Introduction

Architectures based on TR-101[2] related documents have been deployed for about 10 years, allowing operators to promote value-added voice and video services.

In order to broaden support to business and residential customers, encompass fixed and mobile networks, and address both wholesale and retail markets, TR-144 [15] described various requirements including the need for network interconnection standards for broadband access, QoS, bandwidth on demand, increased overall bandwidth, higher network reliability and availability.

Despite a number of significant achievements (e.g. FTTH, IPv6, advanced video services, cloud storage), the end-to-end architecture has not fundamentally changed in recent years. A Residential Gateway (RG) located at the customer premises terminates the home network and connects the LAN devices to the Internet or some other service platforms (e.g., IPTV) through the broadband access network. Over the years, the RG has evolved from a simple modem to a feature rich and complex item. The deployed RG population has become very heterogeneous so it is difficult for a network service provider to obsolete or deprecate a specific RG model. Consequently the deployment of new features or services is often slow and not always possible as older RG models may not have sufficient resources to support a given set of features.

Meanwhile, the NSP's business ecosystem has changed; many of the existing value-added services offered by NSPs (e.g., IPTV, VoIP) have become commoditized while services from Over The Top (OTT) providers are growing at a significant rate. Likewise the consumer electronics industry continues to increase the number and types of devices and services within the residential premises that require broadband connection.

The Information Technology (IT) service provider's ecosystem is changing. Cloud services for residential customers have become commoditized, thanks to the generalization of a number of technologies and techniques used in datacenters (e.g. virtualization, storage). A new generation of technologies and network concepts such as NFV and SDN is emerging.

The purpose of this Technical Report is to address some of the issues foreseen by NSPs for integrating the deployment of new and evolving services into the current architecture, while benefiting from the changing NSP, OTT and IT ecosystems.

The issues that need to be addressed include:

- Reducing the complexity of software management of the RG.
- Reducing the Time to Market (TTM) for deploying a new feature or service that would impact the RG in the current architecture.
- Facilitating the introduction of a Carrier Grade Network Address Translation (CGN) function (in order to alleviate public IPv4 address exhaustion).
- Facilitating troubleshooting of devices and connectivity within the home network

The new use cases that are enabled by the NERG architecture include:

- Providing the NSP with visibility of the devices at the subscriber premises to allow:
    - Downstream QoS to be managed on a per device and/or per user basis

  - Forwarding the subscriber traffic to a centralized parental control platform on a per device and/or per user basis
- A more flexible and agile environment to support a wide set of applications and services: this includes the use of IT technologies and virtualization techniques located in data centers (which have the inherent redundancy/high availability). This also allows technology and commercial scalability ('pay as you grow'). These new applications may be provided by the NSP or third parties, ultimately supporting the concept of an application store for vGs, hence opening the door to new business models.
- Smooth the introduction of M2M services by shifting some of the functions of the M2M home automation devices to the network.
- Offer a new user experience in consuming private or public media content by extending the use of DLNA-like protocols to the WAN, and by shifting some of the functionalities of the Set-Top-Box (STB) to the network.

4.1 NERG Overview

4.1.1 NERG Components

The NERG concept involves moving some of the networking and service-related functions from the RG to the NSP's network. The distribution of functions effectively splits the RG into 2 sets of connected functional components as depicted in Figure 1.

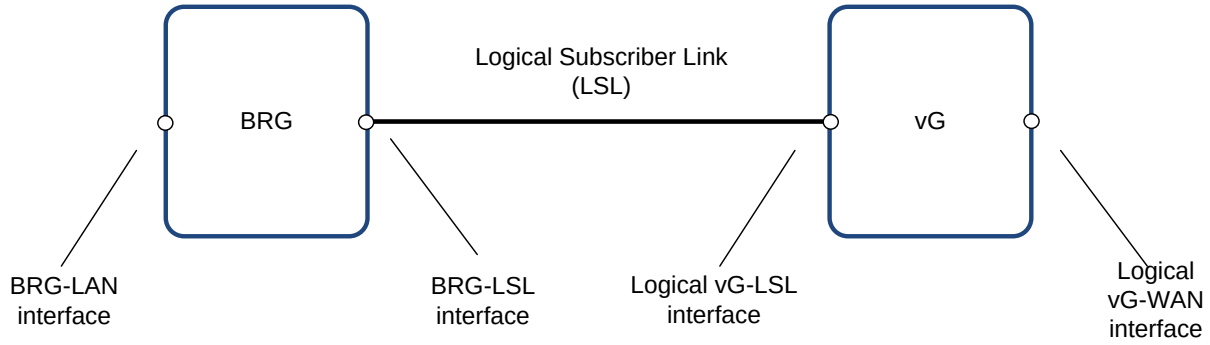

Figure 1 – NERG overview

The main components of the NERG are:

- **BRG**: Bridged Residential Gateway – this is the CPE still located at the residential customer premises, configured as a managed bridge connecting its LAN interfaces and the BRG-LSL interface. This bridge conforms to IEEE 802.1D [3] and 802.1Q [4] and consequently supports MAC learning. Hence local traffic is switched locally, not hair-pinned to the vG.
- **BRG-LAN interface**: Interface(s) on the BRG for connecting LAN devices (i.e. in the home).
- **vG**: virtual Gateway – The set of networking and service related Network Functions that are hosted in the NSP network. At a minimum, the vG provides Network Functions for:
  - the termination of the LSL in the NSP's network,
  - IP address management,

  - IP forwarding and IPv4 NAT capabilities,
  - the termination of the WAN interface

  As the vG provides the capability to terminate the LSL in the NSP's network, the vG is the default IP gateway for the LAN devices.
- **vG-WAN interface**: Logical interface(s) on the vG to one or more IP networks.
- **LSL**: Logical Subscriber Link – the subscriber specific logical point-to-point layer 2 connection between the BRG and the vG. The LSL has the following interfaces:
  - **BRG-LSL interface**: Logical layer 2 interface on the BRG facing the vG.
  - **vG-LSL interface**: Logical layer 2 interface on the vG facing the BRG.

4.1.2 vG Hosting Infrastructure

vGs are hosted on an infrastructure which supports the vG functions (either VNFs or PNFs), terminates the LSLs connecting the BRGs to their respective vGs, and connects the vGs to the IP network. The vG hosting infrastructure may consist of one single item of equipment such as a MS-BNG, or may be composed of the combination of an MS-BNG and some associated NFVI.

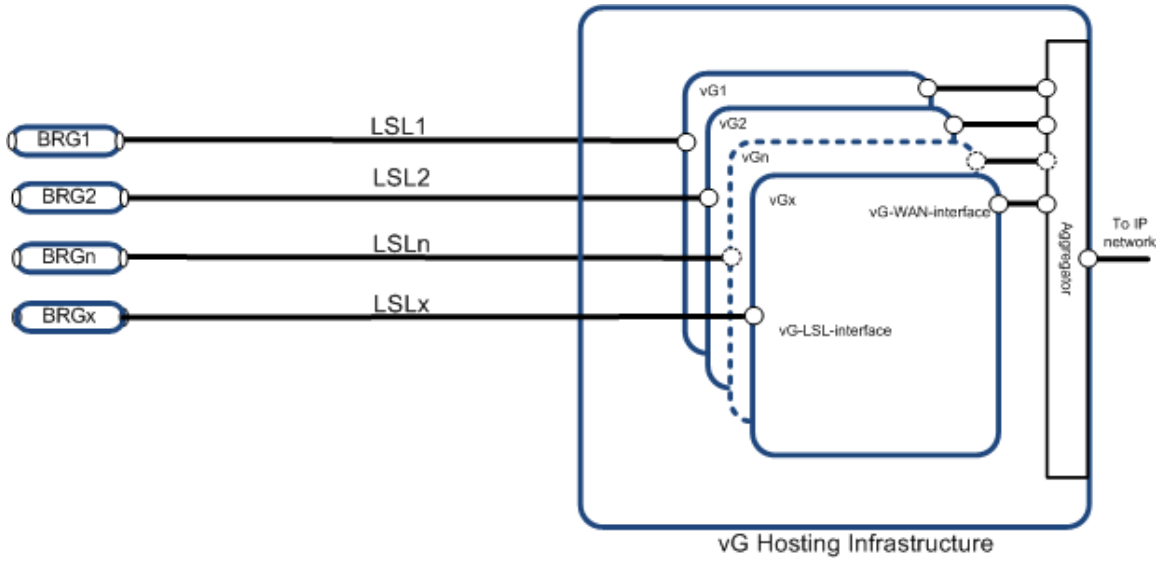

Figure 2 – vG Hosting Infrastructure

4.2 High Level Architectural Components and Functional Distribution Overview

Figure 3 shows the BRG and vG as components of the NERG Architecture over which the Network Functions are distributed. TR-124 [6] defines a set of requirements for RGs in which the following capabilities are identified:

- LAN Interfaces (e.g., Ethernet, WiFi, Voice ATA)
- Multicast
- Learning Bridge
- LAN QoS

- DHCP Client for WAN interface
- DHCP Server for LAN interface
- DNS
- IP WAN Connection
- Upstream QoS
- Downstream QoS
- NA(P)T
- DDoS prevention
- Firewall
- ALG
- Web Re-direct
- WAN Interface
- OAM
- Management Client
- Authentication Client
- PCP Client
- NTP Client
- User Interface

These capabilities are distributed to the BRG or vG depicted in Figure 3 based on the definition of the BRG and vG in section 4.1.

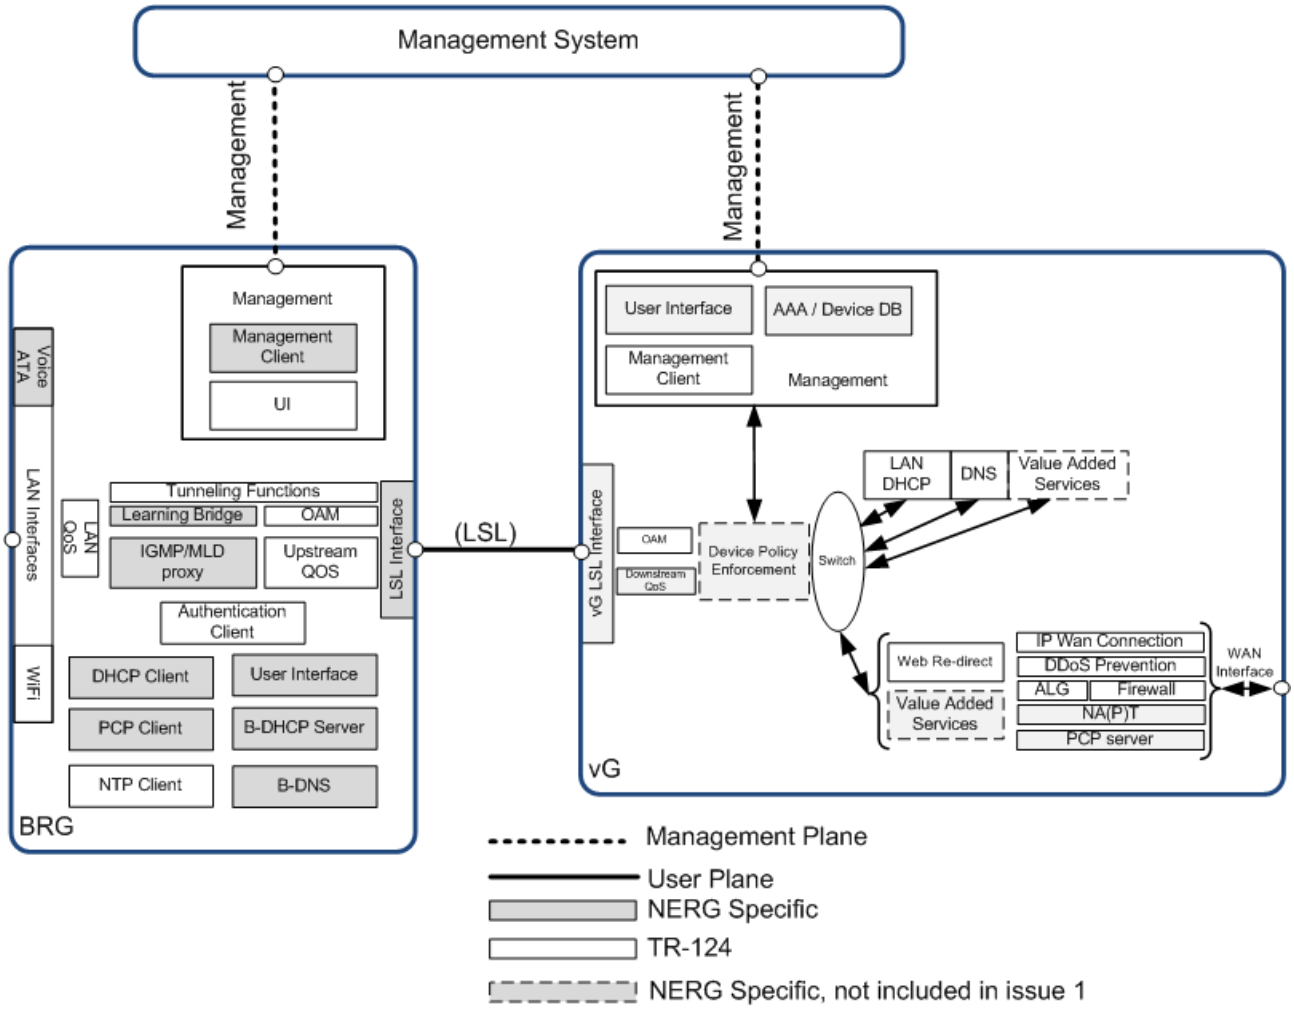

Figure 3 – Functional Distribution of NERG Capabilities

The following functions go beyond TR-124 and are specific to NERG, either to support the connectivity between BRG and vG or to enable NERG specific use cases:

- LSL Interface and tunneling functions – sections 7.1.1 and 7.1.2
- Support for Value Added Services – section 7.5.9
- Device Policy Enforcement – Out of scope of this version

Note: The functional distribution does not prescribe an implementation or system design but is simply to define requirements associated with the identified NERG capabilities. The NERG capabilities can be packaged into Network Functions as needed by the NSP in the BRG and vG, e.g. DHCP, NAT, Parental Control, Firewall, etc.)

4.2.1 BRG functional architecture

The BRG is responsible for transferring data between devices at the customer premises and the vG Functions in the NSP network. The BRG is a managed device and facilitates the troubleshooting of the LSL and devices within the customer premises.

Many of the NERG capabilities which are in the BRG (as depicted in Figure 4) have the same requirements as those defined in TR-124 [6]. However several of the NERG capabilities necessitate additional requirements in the BRG. These include:

- LSL Interface and tunneling functions – section 7.1
- Management Client – section 7.4.4
- DHCP Client – 7.4.3.7
- PCP Client – section 7.4.3.8
- User Interface – section 7.4.6
- Multicast – section 7.4.3.4
- Learning Bridge – section 7.4.3.3
- Voice ATA – section 7.4.2.1.2
- Backup DHCP (B-DHCP) Server – section 7.4.2.3
- Backup DNS (B-DNS) Server – section 7.4.2.4

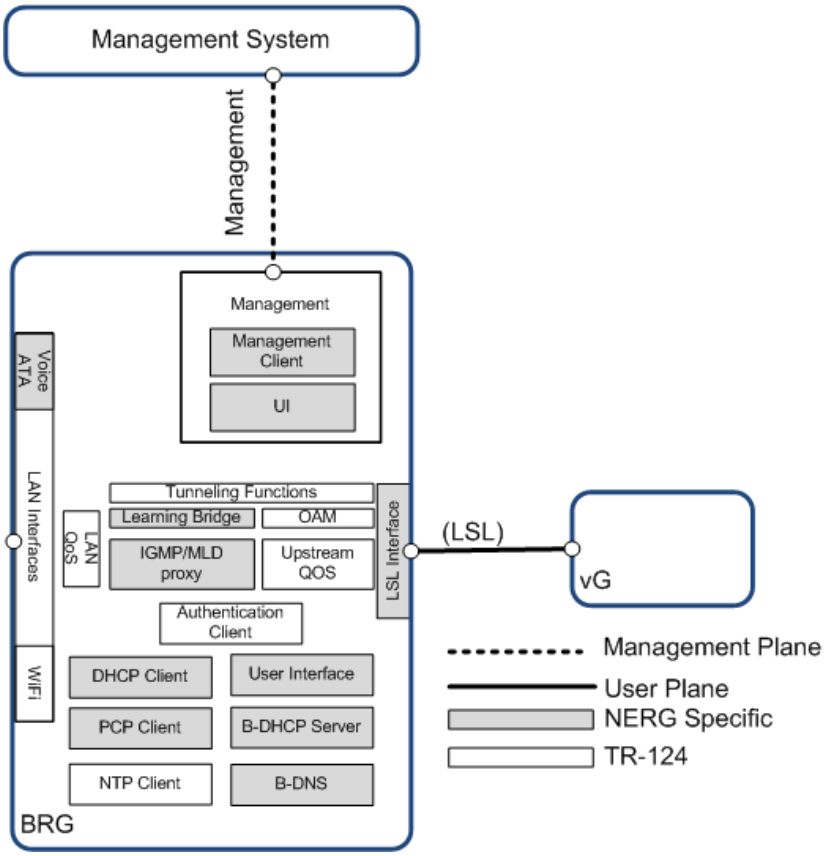

Figure 4 – BRG NERG Capabilities

4.2.2 vG functional architecture

Many of the NERG capabilities moved to the vG as depicted in Figure 5 have the same requirements as those defined in TR-124 [6]. However the following NERG capabilities require additional functionality in the vG:

- LSL Interface – section 7.1
- Extended LAN – section 7.5.10
- LAN DHCP – section 7.5.1
- Device Inventory – Section 7.5.4
- NA(P)T – section 7.5.5
- User Interface – section 7.5.12
- Value Added Services – section 7.5.9
- AAA – section 7.1.3.5

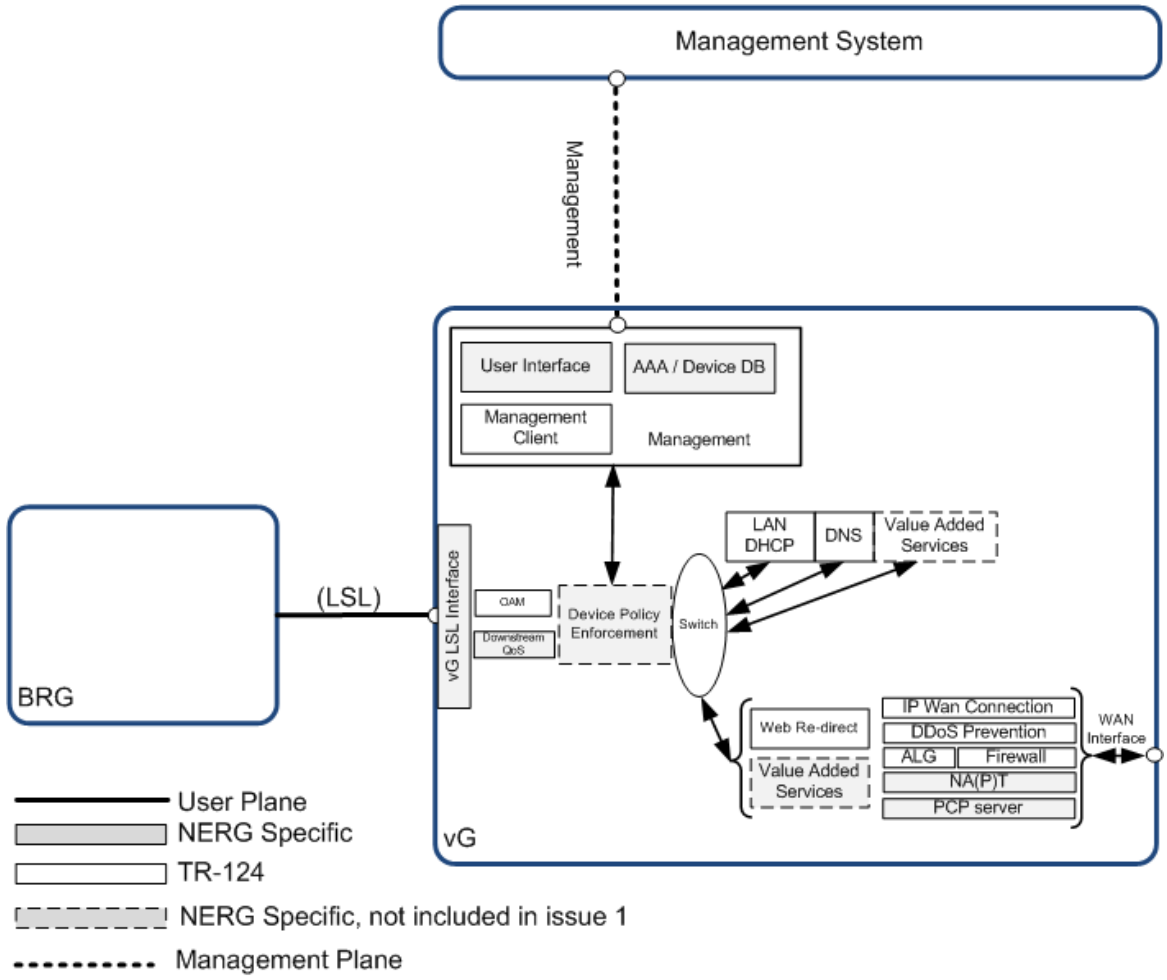

Figure 5 – vG NERG Capabilities

4.2.3 Support for legacy 3-play services

This document only specifies the scenario where services including VoIP, Internet, unicast video and TR-69 management are transported over the LSL, but the managed multicast IPTV service delivered from the access node is NOT transported over the LSL - may or may not have a dedicated VC/VLAN.

This is illustrated by option 1 in the figure below for VoIP service.

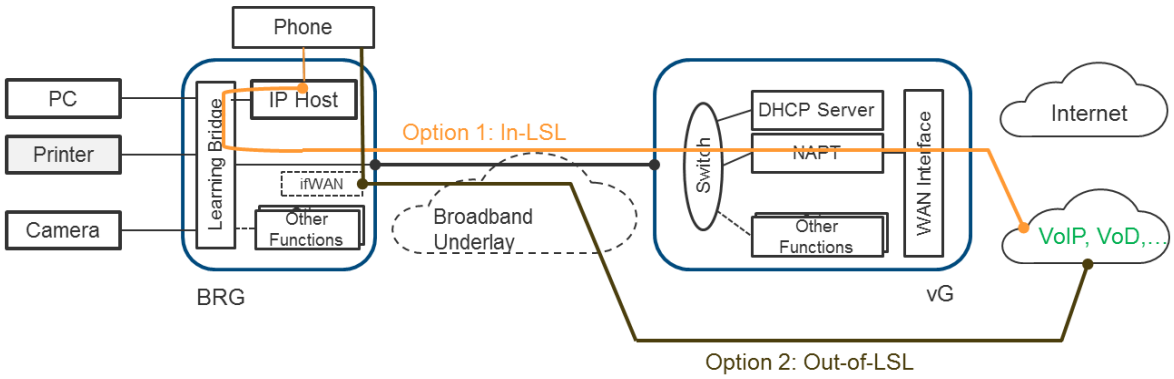

Figure 6 – Support for multi-services

However, other deployment scenarios are possible, although not specified here. For example VoIP can be architected outside of the NERG, using the multi-VC/VLAN architecture as already specified in other documents. This corresponds to option 2 in Figure 6, where VoIP is transported in a dedicated pipe throughout the MSBN using an IP address assigned by the VoIP service platform independently of the LSL and the vG. Similarly, the BRG Management Client is accessed by the remote management system through either the BRG's WAN IP interface or the BRG-LSL interface, but this document only provides specifications for the latter case.

5 NERG Use Cases

Compared to the legacy RG model, the NERG provides additional capabilities including:

- Home device visibility from the vG
- LAN extension to the vG
- Adding network functions to support value added services which do not involve the RG

These enable new use cases, with benefits for both the subscriber and the network service provider.

Note: requirements for per device policies and services are not defined in this issue of the document. However, home device inventory is described in section 7.5.4.

5.1 Device specific services

The vG has MAC address visibility of all the devices on the home LAN; it can therefore categorize devices based on their owner or device profile and apply a specific set of policies. Examples of per device services include:

- Parental control: devices belonging to children can have access to the Internet restricted, for example by time-limits on usage or only being able to access age appropriate content, while parents have unrestricted access.
- Enhanced home office: a work laptop gets bandwidth and prioritization based on a specified QoS policy.
- Assured multimedia: video devices (AppleTV, Chromecast, Roku,…) get the highest QoS priority
- Guest service restrictions: unknown devices can access the Internet, but no other services.

5.2 Security services

The network service provider can offer security services as a value added service (VAS). Examples of security services include firewalling, intrusion detection, web filtering, anti-virus for remote workers, security conscious families and small businesses.

When a subscriber registers for a new security service, a new function is activated in the vG and traffic steering sends all, or a specified subset of, the traffic (e.g. for a given device type) to the service chain that has this security function.

As the security service is hosted in the vG and is activated on demand e.g. via a self-care portal, there is no dependency on changing out the physical RG, and no need for a manual intervention (such as a technician visit or software installation).

5.3 Automatic access to media content

DLNA is a framework for media interoperability between consumer devices within the home network. NERG provides the ability to extend this framework beyond the home network, to resources that are hosted by the network service provider.

The subscriber can access either his own media content hosted in the NSP network (e.g. vNAS), or other media content provided by the operator or some third-party. Virtualized DLNA storage provides a more flexible and scalable solution in terms of storage capacity: the subscriber only pays for what he needs (pay as you grow). In addition, backup and redundancy are now handled by the NSP.

This service leverages NERG's LAN extension and DLNA support in the vG. This can also be combined with improved QoS policies.

5.4 Machine-to-Machine (M2M)

A rapidly growing number of smart connected objects have entered the home, in particular to support smart home applications. These devices are often based on multiple, non-interoperable standards. This creates complexity for the subscriber, who can end up with a plethora of M2M gateways.

The NERG architecture offers a unique opportunity to simplify M2M network design and deployment by introducing an abstraction layer in the operator network that can accommodate existing and future M2M technologies (i.e. different stacks). Being both highly scalable and easily upgradable, this model can accelerate M2M deployment since most the complexity is hidden from a non-technical end-user. The NSP can offer additional services on top of these integrated M2M gateways (e.g. one virtual button to switch off/on the lights in as many rooms as needed, across any M2M technology). The NSP can also facilitate access for third party providers (e.g. utility companies, health care providers) by managing remote access, security, privacy and guaranteeing QoS.

While M2M itself is not specified here, NERG provides key capabilities to allow the NSP to build M2M solutions e.g. hosting of VAS applications in the vG, direct L2 connectivity between the home and the M2M gateway, device visibility, etc..

5.5 IPv6 Migration

In the case where an RG does not support IPv6, NERG can provide an alternative approach to RG replacement by adding IPv6 support in the vG. The existing RG can be changed to a bridge mode configuration (BRG), so that Layer 3, including IPv6 forwarding and control (e.g. DHCPv6 PD, SLAAC), can be delivered by the vG. This may permit IPv6 support without requiring the replacement of the RG.

The vG could also combine IPv6 NERG support with IPv4/IPv6 transition methods, as described in TR-242 Issue 2 [17], as it is deemed to be easier to add these new mechanisms on the vG rather than on physical RGs.

5.6 Enhanced Service Management and Troubleshooting

Operators can take advantage of NERG by allowing all devices connected in the home network to be visible and reachable from the operator's network side (in the traditional RG model, end devices are hidden behind NAT). This direct connectivity can be used to provide enhanced service

management driven by the subscribers' self-care portal as well as optimized troubleshooting by the operator's support technicians.

The subscriber can manage device parameters (IP addressing, filtering, QoS, remote access, etc.) from the portal, which can be accessed from anywhere. Based on the auto discovery of the LAN devices within the customer premises, the support call center is able to check the devices connected to the subscriber's network and launch specific tests as if a support technician was directly connected to the customer network, by extending the customer LAN to the Service Provider's management system for the period of time required.

Reducing support complexity is one of the key benefits expected from NERG, based on the reduced Residential Gateway complexity and the increased visibility and programmability.

6 NERG Management

6.1 Management Functional Architecture

This document does not define the internal architecture of the vG, so only the management of the NERG *functions* is described here. Depending on the vG implementation, other management and orchestration aspects may be necessary (e.g. VNF provisioning and lifecycle management, SDN based service chaining, etc.); this is out of scope for this document.

The NERG functions are distributed between the customer premises and the operator's network which needs to be taken into account from the management perspective. Management of NERG functions includes the configuration, performance monitoring, troubleshooting, and fault management activities associated with NERG functions within the context of a higher layer Service function (e.g., Activation, Diagnostics) implemented by OSS/BSSs.

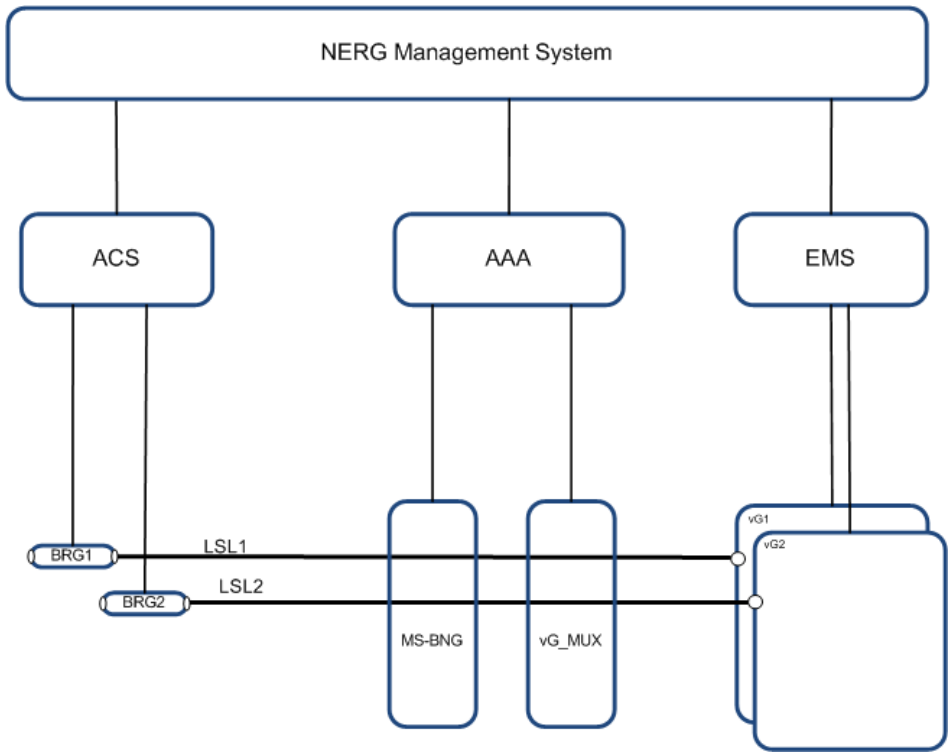

Figure 7 – NERG Management System

Note: if the vG is fully or partially implemented on an NFV infrastructure, the management and orchestration for this NFVI will have to be interfaced with the NERG Management System to manage the vGs.

6.2 NERG User Interface Functional Architecture

The NERG provides user interfaces (UI) to various actors in order to provide direct access to the functions of the NERG. Actors who require access through a User Interface include:

- The Administrator
- End Users to whom the Administrator has authorized access
- A Service Provider Technician
- The Service Provider Help Desk

An actor is typically only authorized to access to a subset of the management functions for a subset of NERG functions. For example, a Service Provider Technician User Interface may only provide access to functions to activate and troubleshoot the NERG.

Actors access the NERG functions through either a User Interface within the Service Provider network or through the BRG's User Interface.

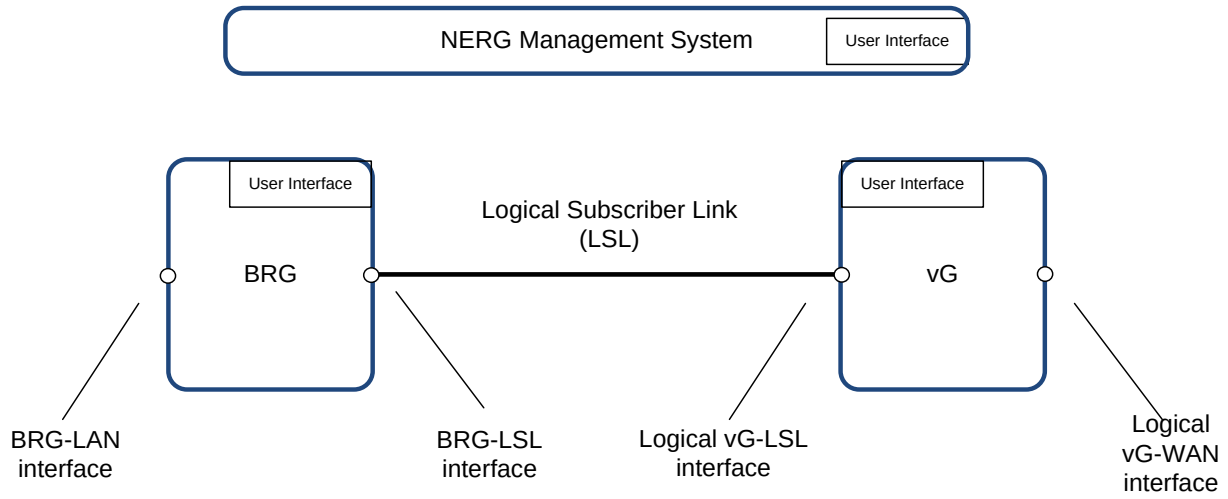

Figure 8 – NERG User Interfaces

Note: There are other interfaces and actors that should not have direct access to the NERG functions but will gain access through the Service Provider's OSS/BSS (e.g. Subscriber for service activation). This type of access is outside the scope of this specification.

6.3 Management Use Cases for NERG

Management of the NERG includes BRG and vG activation, software updates, and troubleshooting of the NERG. This section identifies the main management requirements of the functions of the NERG.

6.3.1 NERG activation

Functions will need to be activated or deactivated within the BRG and/or vG in order for the NERG to support a subscriber's set of services. The NERG management function needs to be able to support the following:

- Download software/firmware and configuration files.

- Activate the functions in the operator's network and/or BRG.
- Accept a notification that the function's host environment is installed and the host environment has applied the software/firmware and artifacts.
- Configure the function including the configuration of the connectivity of the function within the network service.
- Verify the function is operating correctly, for example by checking the operational state and/or executing specific diagnostics.

Once the NERG has had its functions installed and they have been configured, the NERG management function notifies interested parties of the event and that the NERG is operational.

The following sub-sections describe the triggers for the above events.

6.3.1.1 BRG contact activation

The first time a BRG contacts the operator's network, not all of the functions associated with the NERG may be activated on the BRG and within the vG. As such, when the BRG contacts the ACS for the first time, the ACS informs the NERG management system which then may activate additional functions.

6.3.1.2 Software/Firmware updates

Periodically, or as part of a larger service operation (e.g, NERG instance creation), the firmware of the BRG, or software and configuration files associated with one or more functions, might need to be updated. In this scenario, the NERG management function downloads the firmware/software and artifacts into the host environment of the affected functions and performs activities for each subscriber function. Once the affected NERG functions are activated, the NERG management function notifies interested parties of the event and that the NERG is operational.

6.3.1.3 OSS/BSS/User service activation

Services for a subscriber are activated as part of an OSS/BSS subscriber fulfillment process in which the OSS/BSS instructs the NERG management function to activate a set of NERG functions.

6.3.2 In-home connectivity troubleshooting

When managed and non-managed home devices within the customer premises are not functioning properly, subscribers have the ability to attempt to self-diagnose the problem through a self-care portal, or the subscriber is able to contact the operator's help desk for assistance. The troubleshooting interactions needed to remedy the problem with the device may rely on diagnostics that verify the correct operation of one or more functions within the customer premises as well as the operator's network. The OSS interacts with the self-care or help desk portal by sending diagnostic requests for specific functions to the NERG management where the NERG management initiates the requested diagnostic for each function and responds with the result of the each diagnostic request.

Diagnostic requests may be synchronous or asynchronous in nature. If the diagnostic for a function is asynchronous in nature, the NERG management would notify the OSS of the result once the diagnostic has been completed.

6.3.3 Performance monitoring

As part of reactive and proactive monitoring of the NERG by the operator, performance statistics are collected by the NERG which are then reported to OSS systems. Each function might collect performance statistics at a different rate and report the collected statistics on one or more time schedules. The NERG management function is responsible for forwarding the configuration of the collection and reporting function performance statistics to the function being configured.

7 Technical requirements

7.1 E2E Network requirements to support NERG

The vG_MUX (vG Multiplexer) is the entry point of the vG hosting interface. It is responsible for mapping a subscriber to a vG. The LSL that connects a BRG to its vG is composed of two segments: the first segment connects the BRG to the vG_MUX. The second segment extends the LSL from the vG_MUX to the appropriate vG in the vG hosting infrastructure.

The LSL segment from the BRG to vG_MUX can either be:

- a native Ethernet connection, statically or dynamically provisioned through the MSBN. This is the Flat Ethernet LSL connectivity scenario detailed in section 7.1.1.
- an Ethernet connection dynamically established over an IP network using tunneling techniques. This is the Overlay Ethernet LSL connectivity scenario detailed in section 7.1.2.

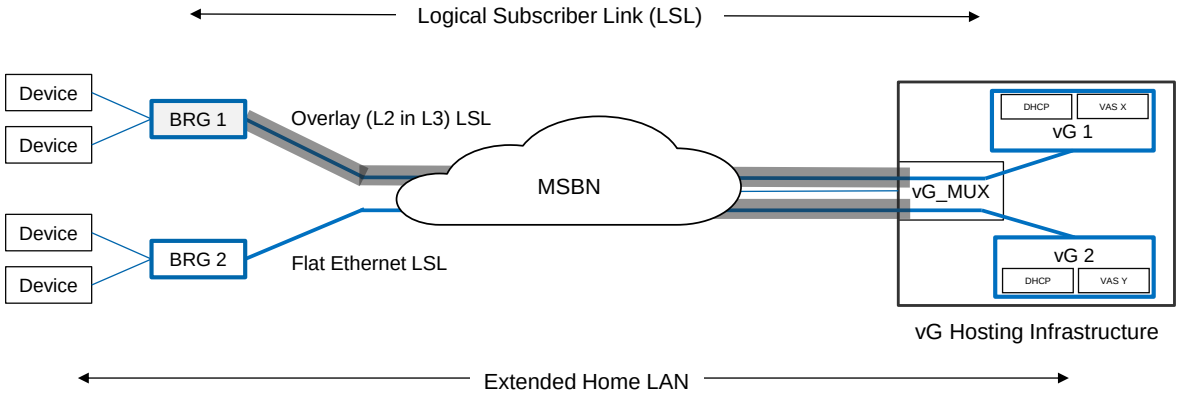

Figure 9 – Flat and Overlay Ethernet architectures

The vG_MUX is a network function that maps L2 traffic between a subscriber's BRG and its unique vG, and ensures traffic isolation between NERG customers. Mapping may be statically provisioned or obtained dynamically via AAA.

The Flat Ethernet LSL scenario is convenient for NSPs who already have a 1-to-1 VLAN access and backhaul architecture in place. It may also allow reusing deployed Residential Gateways by configuring them in bridged mode.

The Overlay LSL scenario simplifies the connectivity to a centralized location such as a data-center because the LSL can be established over a plain IP network. However, some of the existing Residential Gateways may not support tunneling capabilities while meeting throughput requirements.

Note: In the Flat Ethernet LSL case, the BRG-LSL interface and the BRG WAN interface facing the BNG are the same interface. However, some of the traffic which does not originate at or transit via the vG, like multicast (see section 4.2.3) may be carried through this interface but outside of the LSL since the LSL is a logical link between BRG and vG.

Note: PPP support is not considered in this document.

In both Flat and Overlay scenarios, the BRG acts a bridged device, hence the following requirement:

[R-1] The BRG MUST be able to bridge its LAN interfaces and the LSL-WAN-Interface.

7.1.1 Flat Ethernet LSL Connectivity

In the flat model, LSL connectivity can be statically provisioned from the BRG to the vG using the 1:1 VLAN model documented in TR-178 [16] and TR-101[2]. While the BRG to BNG segment is always statically provisioned, the BNG to vG segment can be either statically provisioned, or dynamically constructed by the procedures documented in this section.

Static Model:
The flat model requires the BNG, the vG_MUX and the vG to be provisioned in order to extend the VLAN delineated connectivity between the BRG and the BNG to the vG.

Dynamic Model:
Dynamic provisioning of the LSL segment from the BNG to the vG can be performed at the BNG and the vG_MUX, driven by the AAA operation triggered by the initial DHCP request from the BRG, as illustrated on Figure 10.

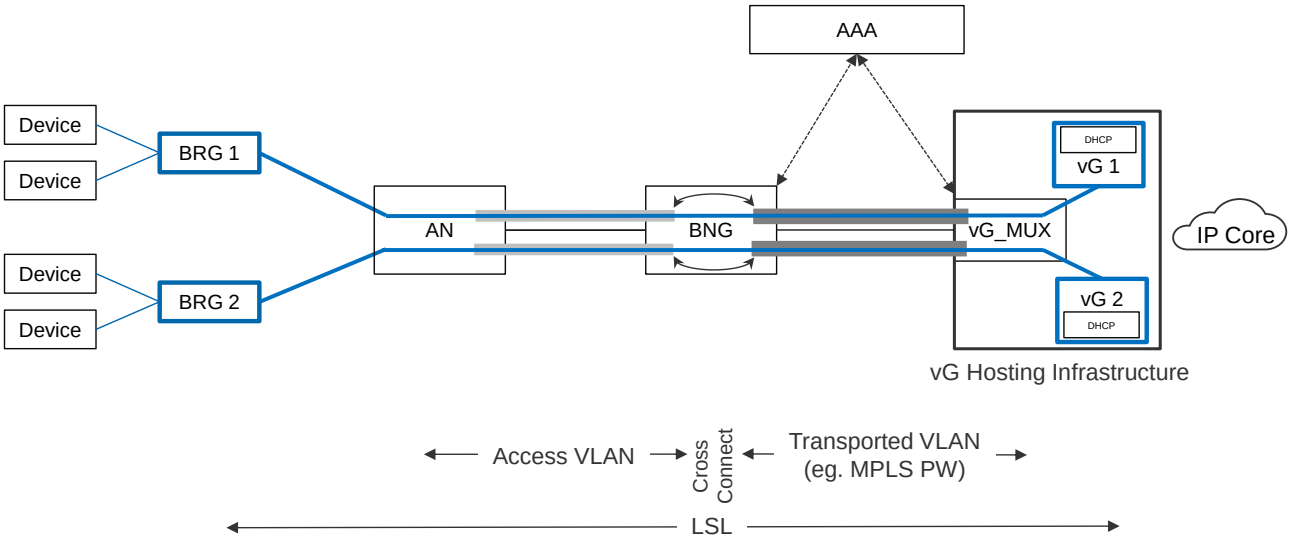

Figure 10 – Flat Ethernet Architecture

The figure below illustrates the call flows starting at BRG boot time and until subscriber devices receive IP addresses.

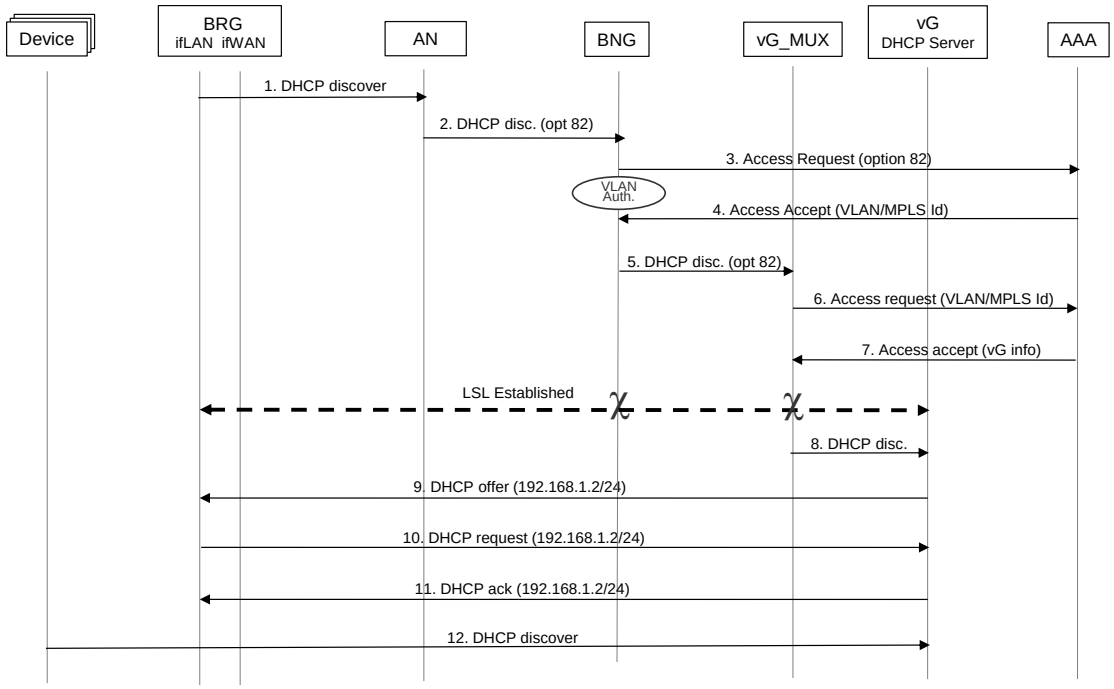

Figure 11 – DHCP and AAA – Stitching of subscriber hosts with vG based on flat model

In the Flat model, the BRG does not have a layer 3 interface on the WAN side. The first DHCP message is issued by the BRG LAN interface and contains option 125 (BBF) with a new sub-option NERG Device Type (24) filled with a string whose value is configurable by the NSP ("BRG" by default). In other words, the BRG only has a private IP address, allocated by the vG.

When the MS-BNG receives a first DHCP request it will buffer the DHCP request, extend the L2 connectivity based on AAA, and relay the DHCP request to the vG. The MS-BNG generates a RADIUS Access-Request to authenticate the customer line. This assumes that the Access Node (AN) is capable of inserting line information in DHCP Option 82, which, in the case of 1:1 VLAN architecture is not mandated in TR-101i2, hence the following requirement:

[R-2] The AN MUST be configurable to act as a DHCP relay as per TR-101i1 section 3.8.2 for 1:1 VLAN tagging mode **M**

The BNG must not override the line identification information, so the following requirement is needed:

[R-3] The BNG implementing section 3.9 of TR-178 MUST be able to be configured to enable/disable the DHCP relay function **M**

The AAA server identifies this is a NERG subscriber and returns to the MS-BNG an Access-Accept message that contains the parameters to extend the VLAN traffic to the node providing the vG_MUX function.

TR-178 [16] provides options to extend L2 traffic from an edge MS-BNG to a centralized MS-BNG. At that point, subscriber L2 traffic is forwarded to the vG_MUX.

Upon receiving a first sign of life from a subscriber line, the vG_MUX maps this L2 traffic to a vG. The mapping between the L2 sub-interface and the targeted vG is dynamically obtained via AAA.

The Access-Request must contain the line identifier information, either option 82 or the dedicated customer identifier between the MS-BNG and vG_MUX, such as VLAN ID or MPLS Identifier. The Access-Accept message returned by the AAA server must contain the forwarding information to get to the appropriate vG as shown in Figure 11. The vG_MUX then generates a RADIUS Acct-Start, to confirm the beginning of the NERG session for this subscriber.

The vG's DHCP server receives this DHCP request from the BRG and assigns an IP address from the subscriber subnet. Other devices connected to the BRG may also request an IP address: at this point, the LSL is already established and none of the previous interactions with AAA on the MS-BNG and vG_MUX are necessary.

7.1.2 Overlay LSL Connectivity

Overlay LSL connectivity is provided by a Layer-2 over IP tunnel between the BRG and the vG. The Overlay Ethernet LSL model works across any TR-178 [16] architecture option, as it uses IP tunnels between the BRG and the vG_MUX function as shown on Figure 12.

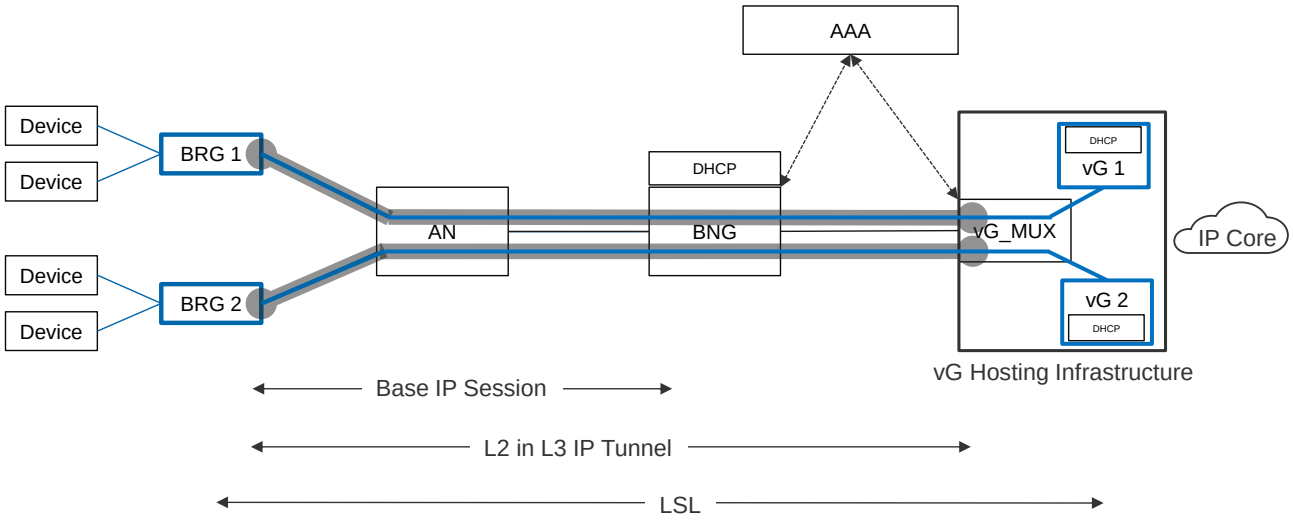

Figure 12 – Overlay LSL architecture

WT-345 [18] will document further network implications if the vG_MUX function (see below) and vGs are hosted in the NFVI.

In general, an overlay approach is agnostic to the network access technology and to the location of the vG infrastructure. For example it allows the NSP to start a field trial or early deployment with a single, highly centralized vG hosting infrastructure, and then decentralize it as the number of NERG subscribers increases.

It also minimizes the impact on the E2E architecture since most of the requirements concern the BRG and the vG themselves. It is also convenient when considering a migration scenario where both legacy and NERG subscribers are supported on the same network infrastructure.

The BRG initiates a Layer-2 over IP tunnel, using the information for tunnel setup (source address, endpoint and protocol) provided by various DHCP options obtained when the initial IP session is established between the BRG and the MS-BNG. On the vG hosting infrastructure, tunnels initiated by the BRGs are terminated by the vG_MUX. This function extracts (decapsulates) the upstream Ethernet frames and forwards them to the appropriate vG based on the BRG WAN interface IP address. Likewise, this function steers downstream Ethernet frames originating from the vG to the appropriate tunnel.

The figure below illustrates the call flows at BRG boot time.

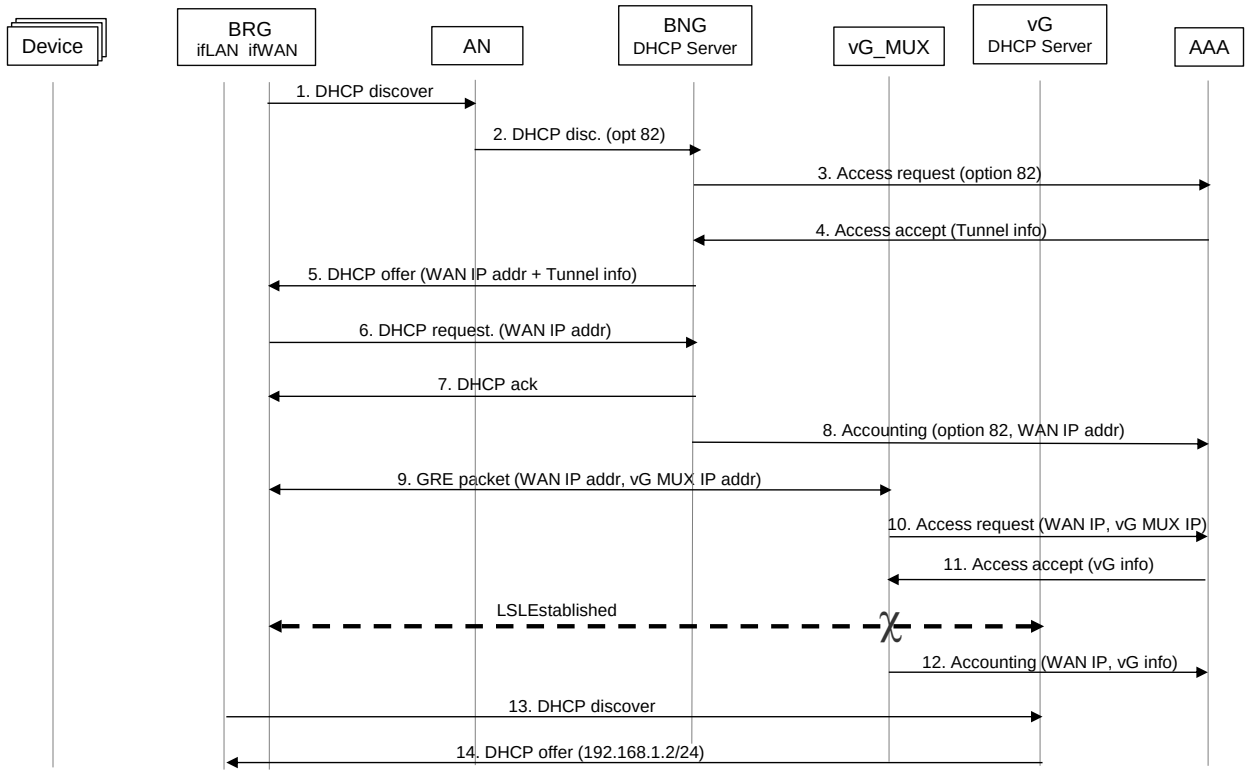

Figure 13 – DHCP and AAA – Tunneling information statically configured on DHCP server

At boot time, the BRG requests an IP address for its WAN interface with a DHCP request message containing option 125 (BBF) with new sub-option NERG Device Type (24) filled with the string whose value (“BRG” by default) is configurable by the NSP. This option combined with option 82 is used to let the AAA and DHCP server know that the device requesting an IP address is a BRG. The MS-BNG intercepts the BRG DHCP request and generates a RADIUS Access-Request to authenticate the BRG. Authentication requires options to ensure that the BRG can be uniquely identified. This can be done using the line ID inserted by the access node (option 82).

The DHCP server is either an independent platform (the MS-BNG simply acting as a DHCP relay) or hosted by the MS-BNG. The required tunneling information is either configured in the DHCP server or provided by the AAA infrastructure as shown on Figure 13. If the tunneling information is provided through AAA, the RADIUS attributes to be used and how they are translated into DHCP options is detailed in section 7.1.3.3.2.

The DHCP server assigns an IP address to the BRG and also provides the tunneling information through newly defined DHCP options (tunnel endpoint, tunneling protocols as detailed in section 7.1.3.2).

The BRG then needs to obtain an IP address for its LAN interface. It sends a DHCP request over the tunnel to the vG's DHCP server. When vG_MUX receives the first packet encapsulated in the L2-tunnel, it needs to know which vG it has to forward it to. The mapping between the tunnel source IP address which identifies the subscriber and the targeted vG can be either statically configured within the vG_MUX function, or dynamically obtained via AAA.

In the case of AAA, vG_MUX sends a RADIUS access-request to the AAA server to find the mapping between subscriber source address and corresponding vG. Consequently, the AAA server must encode in its response the transport protocol and the set of parameters specific to this protocol, using the RADIUS attributes listed in section 7.1.3.4.

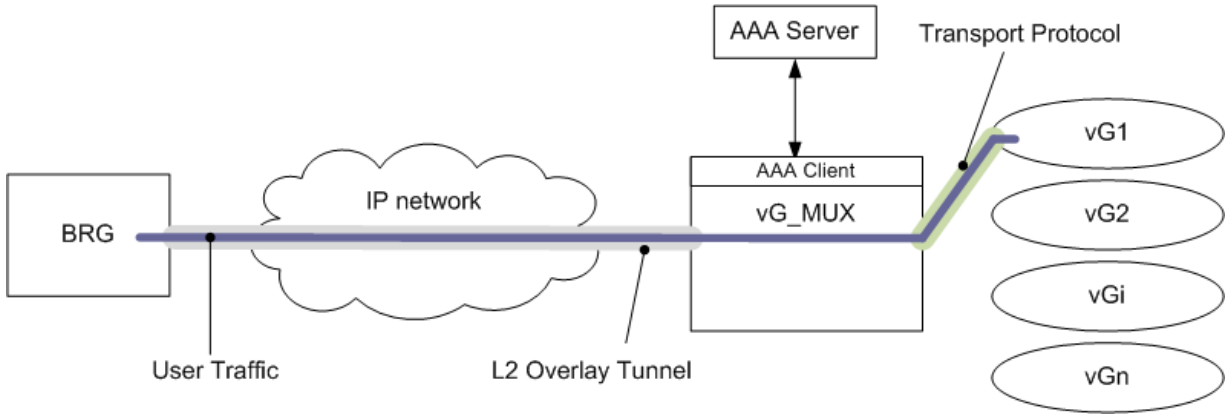

Figure 14 – vG_MUX function and connection to vG

The vG_MUX decapsulates the tunneled traffic and forwards it to the now identified and reachable vG. The vG_MUX then generates a RADIUS Acct-Start as the beginning of the NERG session for this subscriber. The vG's DHCP server receives this DHCP request from the BRG and assigns an IPv4 address to its LAN interface. This address belongs to the subscriber's private IP subnet, e.g. 192.168.1.0/24.

7.1.2.1 Support for Tunneling – Requirements on the BRG

To setup the Overlay LSL, the BRG must be provided with the necessary tunneling information via DHCP. Tunneling information can also be provided via TR-69 outside the LSL, but this is out of scope of this document. Consequently the BRG must meet the following set of requirements:

[R-4] The BRG MUST be able to obtain the IP configuration of its network interface, through DHCP, prior to tunnel establishment.
[R-5] The BRG MUST be able to establish a L2 tunnel over IP to vG_MUX using the information received via DHCP.
[R-6] The BRG MUST support Ethernet over GRE tunneling

7.1.2.2 Support for Tunneling – Requirements on the DHCP server

The DHCP server must be able to differentiate between NERG and legacy RGs on the basis of information received from the AAA server. The DHCP server assigns an IPv4 or an IPv6 address, and in the case of BRG, also provides the BRG with tunneling information obtained from the AAA server and translated into the DHCP options specified in section 7.1.3.3.2.

[R-7] The DHCP server MUST be able to translate RADIUS tunnel attributes into DHCP option 125 encoding of the LSL tunnel information using BBF sub-options.

7.1.2.3 Support for Tunneling – Requirements on the vG hosting infrastructure

The vG hosting infrastructure must be able to terminate the tunnel initiated by the BRG and associate its Ethernet traffic with the vG LSL-interface of a vG. This requires mapping a tunnel instance to a vG, so as to associate the BRG with the vG.

[R-8] The vG hosting infrastructure MUST support a vG_MUX function to terminate L2 over IP tunnels.
[R-9] The vG_MUX MUST transfer Ethernet traffic between the tunnel endpoint and its associated vG.
[R-10] The vG_MUX MUST support the Ethernet over GRE tunneling protocol as per RFC-2784 [8].
[R-11] The LSL tunnel MUST uniquely bind the home LAN to the vG at the vG_MUX function. **M C**

Note: The uniqueness can be based on tunnel endpoints (tuple of BRG WAN Interface IP address and vG_MUX IP address).

7.1.2.4 MTU considerations

The network nodes in access and backhaul networks need to support an MTU of at least 1518 bytes to transport L2 over IP without causing fragmentation.

[R-12] The MTU at the BRG MUST be configurable. **M**
[R-13] The BRG MUST support an MTU that permits Ethernet frames of at least 1518 bytes to be encapsulated within the LSL tunnel.
[R-14] The MTU at the vG_MUX MUST be configurable. **M**
[R-15] The vG_MUX MUST support an MTU that permits Ethernet frames of at least 1518 bytes to be encapsulated within the LSL tunnel.

7.1.3 AAA Requirements

AAA is used in NERG architecture at different nodes and for different purposes, as shown on Figure 15:

- At the MS-BNG:
  - in the case of the Flat Ethernet LSL, to extend the subscriber's access VLAN to the vG_MUX, by dynamically cross-connecting the Access-VLAN to an L2 network resource such as a VxLAN, an MPLS pseudo-wire or an Ethernet over GRE tunnel (case #1 on Figure 15),
  - in the case of the Overlay Ethernet LSL, to provide the BRG with the necessary tunneling information via DHCP to allow the BRG to connect to vG_MUX. In this case, the tunneling information is encoded in a set of new RADIUS Attributes and translated into a set of DHCP options (case #2). In this document, only Ethernet over GRE tunneling is defined.
- At the vG_MUX:
  - to extend the LSL from the vG_MUX to the subscriber's vG, the LSL being either a Flat Ethernet LSL (case #3) or an Overlay LSL (case #4).

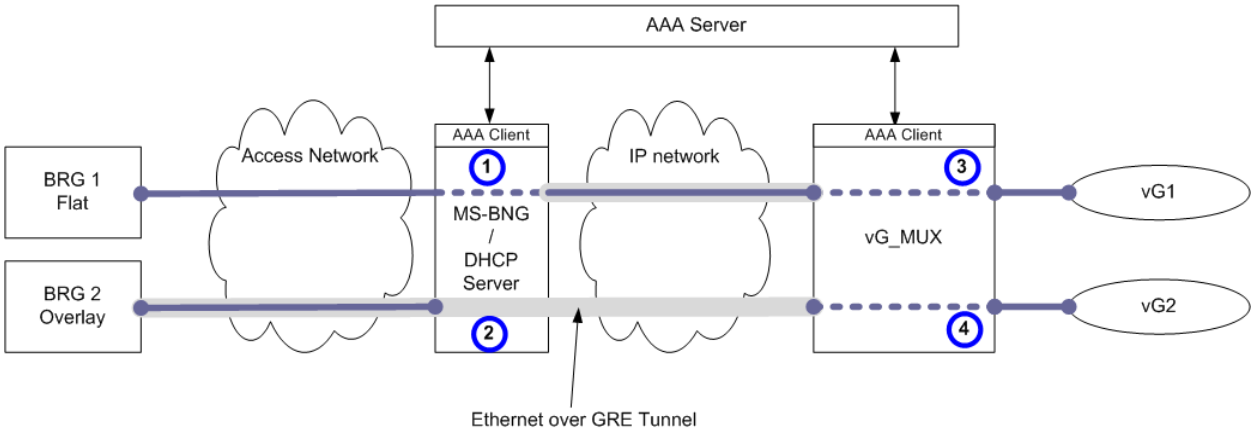

Figure 15 – AAA at the MS-BNG and at vG_MUX

7.1.3.1 AAA at the MS-BNG

In order to create connectivity from the BRG to the vG_MUX, the MS-BNG needs to be able to:

- extend the subscriber's access VLAN to the vG_MUX, in case of Flat Ethernet LSL
- provide the BRG with tunneling information, in the case of Overlay Ethernet LSL.

This requires the below new RADIUS attributes, defined via Broadband Forum Vendor Specific Attributes (VSA):

BBF-LSL-Tunnel-Type:
Description: This attribute indicates the tunnel type to be used for the LSL.

- When the tunnel-type is FLAT, the MS-BNG will use the AAA information to perform the local LSL extension configuration.

- When the tunnel-type is not FLAT, the MS-BNG will use DHCP to pass the relevant information to the BRG.

Values:

- 0: SoftGRE (Ethernet over GRE)
- 1: VxLAN (defined value, but support for VxLAN tunneling is out scope of issue 1)
- 2: L2TPv3 (defined value, but support for L2TPv3 tunneling is out of scope of issue 1)
- 3: FLAT

BBF-LSL-Server-Endpoint-v4:
Description: This attribute defines the Ethernet over GRE tunnel end-point at the vG_MUX side.
Value: IPv4 address

BBF-LSL-Server-Endpoint-v6:
Description: This attribute defines the Ethernet over GRE tunnel end-point at the vG_MUX side.
Value: IPv6 address

BBF-LSL-Server-Endpoint-FQDN:
Description: This attribute defines the Ethernet over GRE tunnel end-point at the vG_MUX side.
Value: FQDN

Note: 3 different RADIUS attributes are used for the endpoint addresses, depending if it is IPv4, IPv6 or FQDN. This is to respect RFC 6158, RADIUS Design Guidelines, which discourages using polymorphic attributes.

BBF-LSL-Client-Endpoint-v4:
Description: This attribute defines the IPv4 address that the BRG needs to use to set up its tunnel. This attribute may be needed in the case where the BRG has multiple IP addresses on its WAN interfaces (for example, when the VoIP service is supported in a dedicated VLAN, outside the LSL)
Value: IPv4 address.

BBF-LSL-Client-Endpoint-v6:
Description: This attribute defines the IPv6 address that the BRG needs to use to set up its tunnel. This attribute may be needed in the case where the BRG has multiple IP addresses on its WAN interfaces (for example, when the VoIP service is supported in a dedicated VLAN, outside the LSL)
Value: IPv6 address.

BBF-LSL-Client-Endpoint-FQDN:
Description: This attribute defines the FQDN pointing to the IPv4 or IPv6 address that the MS-BNG or vG_MUX needs to use for setting up its tunnel.
Value: FQDN

The new RADIUS attribute below is required both on MS-BNG (flat LSL case) and on the vG_MUX :

BBF-LSL-Tunnel-Private-Group-ID

Description: Specifies the ID of the network resource to use at the MS-BNG or at vG_MUX to extend the LSL respectively to vG_MUX or to the vG.
Value: string encoding the ID of the network resource (VLAN ID, VNI, MPLS label, etc.).

7.1.3.2 AAA at the MS-BNG – Flat LSL (case "1")

In the Flat LSL scenario, the MS-BNG cross-connects the Subscriber's Access-VLAN to an L2 connection ending at the vG_MUX. This connection needs to be pre-provisioned by the NSP. This connection may be a VLAN, a VxLAN or an MPLS pseudo-wire (VPWS).

7.1.3.2.1 Extension to vG_MUX using a VLAN

RFC 3580 section 3.31 [11] defines the Tunnel-Type VLAN and provides a method using 3 RADIUS attributes to dynamically place a subscriber port into a particular VLAN, based on the result of the authentication:

```
BBF-LSL-Tunnel-Type=FLAT (3) # to tell the MS-BNG to do FLAT actions.
Tunnel-Type=VLAN (13)# VLAN (RFC 3580)
Tunnel-Medium-Type=802 (6) # Ethernet (RFC 2868)
Tunnel-Private-Group-ID=VLANID # VLAN ID encoded as a string (RFC 2868).
```

[R-16] For Flat Ethernet LSL, the MS-BNG MUST be able to map the BRG's traffic into a VLAN using the above set of attributes

Note: depending of the implementation, additional attributes may be required (e.g. egress port, stacked VLAN, etc.).

7.1.3.2.2 Extension to vG_MUX using a VxLAN

VxLAN is currently not a Tunnel-Type defined in RFC 2868 [9].

This requires a new Tunnel-Type to be defined by the IETF. The following set of attributes could then be used to map a subscriber port into a particular VxLAN:

```
BBF-LSL-Tunnel-Type=FLAT (3) # to tell the MS-BNG this is FLAT mode.
Tunnel-Type=VxLAN(tbd # tunnel-type value for VxLAN
BBF-LSL-Tunnel-Server-Endpoint-IPv4=IPv4 address # VTEP on vG_MUX.
BBF-LSL-Tunnel-Server-Endpoint-IPv6=IPv6 address # VTEP on vG_MUX.
BBF-LSL-Tunnel-Server-Endpoint-FQDN=FQDN # VTEP on vG_MUX.
BBF-LSL-Tunnel-Client-Endpoint-IPv4=IPv4 address # VTEP on MS-BNG
BBF-LSL-Tunnel-Client-Endpoint-IPv6=IPv6 address # VTEP on MS-BNG
BBF-LSL-Tunnel-Client-Endpoint-FQDN=FQDN to IP address # VTEP on MS-BNG
BBF-LSL-Tunnel-Private-Group-ID=VNI # VNI encoded as a string.
```

[R-17] For Flat Ethernet LSL, the MS-BNG MUST be able to map the BRG's traffic into a VxLAN using the above set of attributes.

7.1.3.2.3 Extension to vG_MUX using an MPLS-VPWS

This requires a new Tunnel-Type to be defined. The following set of attributes can then be used to map a subscriber port into a particular PW VPWS:

```
BBF-LSL-Tunnel-Type=FLAT (3) # to tell the MS-BNG to do FLAT actions.
Tunnel-Type=VPWS(tbd))# tunnel-type value for VPWS
BBF-LSL-Tunnel-Server-Endpoint-IPv4=IPv4 address # PW endpoint address on vG_MUX
BBF-LSL-Tunnel-Server-Endpoint-IPv6=IPv6 address # PW endpoint IPv6 address on vG_MUX
BBF-LSL-Tunnel-Server-Endpoint-FQDN=FQDN # PW endpoint FQDN on vG_MUX.
BBF-LSL-Tunnel-Client-Endpoint-IPv4=IPv4 address # PW endpoint address on MS-BNG
BBF-LSL-Tunnel-Client-Endpoint-IPv6=IPv6 address # PW endpoint IPv6 address on MS-BNG
BBF-LSL-Tunnel-Client-Endpoint-FQDN=FQDN to IP address # PW endpoint FQDN on MS-BNG
BBF-LSL-Private-Group-ID=LABEL # PW Label encoded as a string
```

[R-18] For Flat Ethernet LSL, the MS-BNG MUST be able to map the BRG's traffic into a VPWS PW using the above set of attributes.

7.1.3.3 AAA at the MS-BNG – Overlay LSL (case "2")

7.1.3.3.1 Overlay LSL using Ethernet over GRE Tunneling

In the case of Overlay LSL, this document only specifies the use of Ethernet over GRE. The AAA server is used to indicate to the MS-BNG that the Overlay technique is being used, and to provide the MS-BNG and/or the DHCP server with tunneling information. The AAA server would typically send an Access-Accept message containing the following attributes:

```
Framed-IP-Address=IPv4 address #assigned to BRG's WAN interface
BBF-LSL-Tunnel-Type: 0 (GRE) # Overlay Ethernet over GRE.
BBF-LSL-Tunnel-Server-Endpoint-IPv4=IPv4 address # endpoint at vG_MUX
BBF-LSL-Tunnel-Server-Endpoint-IPv6=IPv6 address # endpoint at vG_MUX
BBF-LSL-Tunnel-Server-Endpoint-FQDN=FQDN # endpoint at vG_MUX
BBF-LSL-Tunnel-Client-Endpoint-IPv4=IPv4 address # endpoint at BRG
BBF-LSL-Tunnel-Client-Endpoint-IPv6=IPv6 address # endpoint at BRG
```

7.1.3.3.2 Translation of GRE RADIUS Attributes into DHCP and DHCPv6 options

If the DHCP server obtains overlay tunneling information through AAA, it needs to translate these attributes and their values into DHCP or DHCPv6 options to provide this tunneling information to the BRG.

If BBF-LSL-Tunnel-Server-FQDN is used, the DHCP server resolves the FQDN and encodes the resulting IPv4 or IPv6 address as per the following requirements:

[R-19] The DHCP server MUST be able to map attribute BBF-LSL-Tunnel-Type to DHCPv4 option 125 sub-option 21.

[R-20] The DHCP server MUST be able to map attribute BBF-LSL-Server-Endpoint-IPv4 to DHCPv4 option 125 sub-option 22.

[R-21] When BBF-LSL-Server-Endpoint-FQDN is contained in the RADIUS response, the DHCP server MUST resolve the DNS FQDN. If the result is an IPv4 address, the DHCP server MUST be able to map this IPv4 address to DHCPv4 option 125 sub-option 22.

[R-22] The DHCP server MUST be able to map attribute BBF-LSL-Client-Endpoint-IPv4 to DHCPv4 option 125 sub-option 23.

[R-23] When BBF-LSL-Client-Endpoint-FQDN is contained in the RADIUS response, the DHCP server MUST resolve the DNS FQDN. If the result is an IPv4 address, the DHCP server MUST be able to map this IPv4 address to DHCPv4 option 125 sub-option 23.

[R-24] The DHCP server MUST be able to map attribute BBF-LSL-Server-Endpoint-IPv6 to DHCPv6 option 17 sub-option 22.

[R-25] When BBF-LSL-Server-Endpoint-FQDN is contained in the RADIUS response, the DHCP server MUST resolve the DNS FQDN. If the result is an IPv6 address, the DHCP server MUST be able to map this IPv6 address to DHCPv6 option 17 sub-option 22.

[R-26] The DHCP server MUST be able to map attribute BBF-LSL-Client-Endpoint-IPv6 to DHCPv6 option 17 sub-option 23.

[R-27] When BBF-LSL-Client-Endpoint-FQDN is contained in the RADIUS response, the DHCP server MUST resolve the DNS FQDN. If the result is an IPv6 address, the DHCP server MUST be able to map this IPv6 address to DHCPv6 option 17 sub-option 23.

7.1.3.4 AAA at the vG_MUX (Cases “3” & “4”)

The vG_MUX function, on the vG hosting infrastructure side, needs to know to which vG the user traffic coming from a given BRG must be forwarded. When it receives the first packet from a BRG (containing the VLAN tag in Flat model or the tunnel source IP address in the Overlay model), the vG_MUX sends a RADIUS request to the AAA server. The AAA server must indicate to the vG_MUX how to forward the subscriber’s traffic to its vG.

7.1.3.4.1 VLAN transport

As per the mechanism introduced in 7.1.3.2.1, the following attributes are used:

```
Tunnel-Type=VLAN (13)# VLAN (RFC 3580)
Tunnel-Medium-Type=802 (6)# Ethernet (RFC 2868)
Tunnel-Private-Group-ID=VLANID # VLAN ID encoded as a string (RFC 2868)
```

[R-28] The vG_MUX function MUST be able to map the BRG’s Ethernet traffic into a VLAN using the above set of attributes.

Note: depending of the implementation, additional attributes may be required (e.g. egress port, stacked VLAN, etc.).

7.1.3.4.2 VxLAN Transport

This requires a new Tunnel-Type to be defined. The following set of attributes can then be used to map a subscriber port into a particular VxLAN:

```
BBF-LSL-Tunnel-Type=VxLAN(tbd) # tunnel-type value for VxLAN
BBF-LSL-Tunnel-Server-Endpoint-IPv4=IPv4 # VTEP IPv4 at vG
BBF-LSL-Tunnel-Server-Endpoint-IPv6=IPv6 # VTEP IPv6 at vG
```



```
BBF-LSL-Tunnel-Client-Endpoint-IPv4=IPv4 # VTEP IPv4 on vG_MUX
BBF-LSL-Tunnel-Client-Endpoint-IPv6=IPv6 # VTEP IPv6 on vG_MUX
BBF-LSL-Tunnel-Private-Group-ID=VNI # VNI encoded as a string
```

[R-29] The vG_MUX function MUST be able to map the BRG's tunneled traffic into a VxLAN using the above set of attributes.

7.1.3.4.3 MPLS-VPWS Transport

This requires a new Tunnel-Type to be defined. The following set of attributes can then be used to map a subscriber port into a particular PW:

```
BBF-LSL-Tunnel-Type=VPWS (tbd)# tunnel-type value for MPLS-VPWS
BBF-LSL-Tunnel-Server-Endpoint-IPv4=IPv4 address # PW endpoint address on vG
BBF-LSL-Tunnel-Server-Endpoint-IPv6=IPv6 address # PW endpoint IPv6 address on vG
BBF-LSL-Tunnel-Server-Endpoint-FQDN=FQDN # PW endpoint FQDN on vG
BBF-LSL-Tunnel-Client-Endpoint-IPv4=IPv4 address # PW endpoint address on vG_MUX
BBF-LSL-Tunnel-Client-Endpoint-IPv6=IPv6 address # PW endpoint IPv6 address on vG_MUX
BBF-LSL-Tunnel-Client-Endpoint-FQDN=FQDN to IP address # PW endpoint FQDN on vG_MUX
BBF-LSL-Private-Group-ID=LABEL# PW Label encoded as a string.
```

[R-30] The vG_MUX function MUST be able to map the BRG's tunneled traffic into a VPWS PW using the above set of attributes.

7.1.3.5 AAA for device identification and profile enforcement

Home device identification for specific device policy enforcement is out of scope of this document. However, the vG could intercept the devices' DHCP messages and send RADIUS requests accordingly to a AAA infrastructure dedicated to home devices, so to enforce specific QoS policy and/or service chaining for example. In this case, the vG itself is a AAA client.

7.1.4 Home network Availability

7.1.4.1 Introduction

The NERG architecture requires all the elements in the path between the BRG and the vG through the LSL as well as the vG's WAN connection to be operational, since a set of vital functions are shifted to the vG. Partial or complete failure of any of the elements in this path would result in partial or complete failure of services, and could lead to connectivity issues within the home network. In addition, the complexity of the vG hosting infrastructure where the vG is thosted in one or more network elements or hosts, potentially increases the chance of a failure.

These connectivity issues would impact the customer experience by not only losing access to the Internet, but also preventing capabilities such as:

- printing a document on a printer in the home (LAN) network.
- running M2M applications in the home,
- playing content on a media server in the home network.

When a failure occurs either in the vG, in the vG hosting infrastructure or on the LSL that would cause connectivity issues between devices in the home network, the BRG can provide a temporary substitute for the necessary network-based functionality. The following functions have been identified as being needed in this situation:

- A backup DHCP server (B-DHCP): enables devices within the home network to obtain or renew their IP address.
- A backup DNS server (B-DNS): enables access to devices within the home network via a domain name.

7.1.4.2 LSL Failure Detection

LSL failures are detected by the BRG monitoring communication with the vG as specified in section 7.1.5.

7.1.4.3 B-DHCP Server Activation

When the LSL is operational, the BRG needs to snoop the DHCP exchanges to maintain a database of current IP address leases.

When an LSL failure is detected, this database permits the role of the DHCP server to be taken over by the B-DHCP Server.

The B-DHCP server is able to respond to DHCP RENEW messages sent by devices which already have an IP address using the information in the database. When a device has no database entry, the B-DHCP server assigns a new IP address with a short lease as specified in section 7.4.2.3.

[R-31] The BRG MUST support a B-DHCP server.
[R-32] The BRG MUST snoop the DHCP messages between the devices and the vG DHCP server .
[R-33] The BRG MUST maintain a lease database.
[R-34] The B-DHCP server MUST be activated on LSL failure detection.

7.1.4.4 Restoration of the vG DHCP Server

Once the LSL becomes available again, DHCP requests are served by the vG DHCP server. Consequently, the B-DHCP server needs to be deactivated.

[R-35] When the LSL becomes available again, the B-DHCP server MUST NOT respond to any DHCP requests.
[R-36] When the LSL becomes available again, the B-DHCP MUST issue a FORCERENEW message to devices to which it has assigned or renewed an address.

Note: Because the vG DHCP server is not aware of the addresses assigned by the B-DHCP server during the failure, there is a risk of IP address collision: at the same time, one device D1 has an IP address that was assigned by B-DHCP server during the failure and another device D2 has the same IP address that was assigned by the vG DHCP server after the reestablishment, before D1's lease expires.

This risk is limited by the short leases used by the B-DHCP server.
However, in order to completely mitigate the risk, subnetting of the address pool could be used, so that the B-DHCP and DHCP servers will not use overlapping address ranges. The specifics of subnetting are out of scope for this document.

7.1.4.5 B-DNS Server Activation, and Restoration of the vG DNS Server

When the DNS Server in the vG becomes unavailable, the BRG activates a B-DNS Server that provides a name service for devices in the home network (e.g.: user-friendly printer names, user interface access).

[R-37] When the DNS in the vG is unavailable, the BRG MUST activate its B-DNS .

[R-38] When the DNS in the vG becomes available, the BRG MUST de-activate its B-DNS.

[R-39] When the B-DNS Server is deactivated, the B-DNS Server MUST NOT respond to any DNS requests.

7.1.5 Connectivity Management & LSL monitoring

Both ends of the LSL need to monitor the availability of connectivity between the BRG and the vG. The BRG uses this information to control the operation of the backup DHCP server (as per section 7.1.4.3). The vG may use this information for resource management. The monitoring mechanism is based on ARP requests as described in TR-146 [20] for IP session monitoring and ICMP ping.

Figure 16 below shows the ARP ping keep-alive mechanism on both BRG and vG.

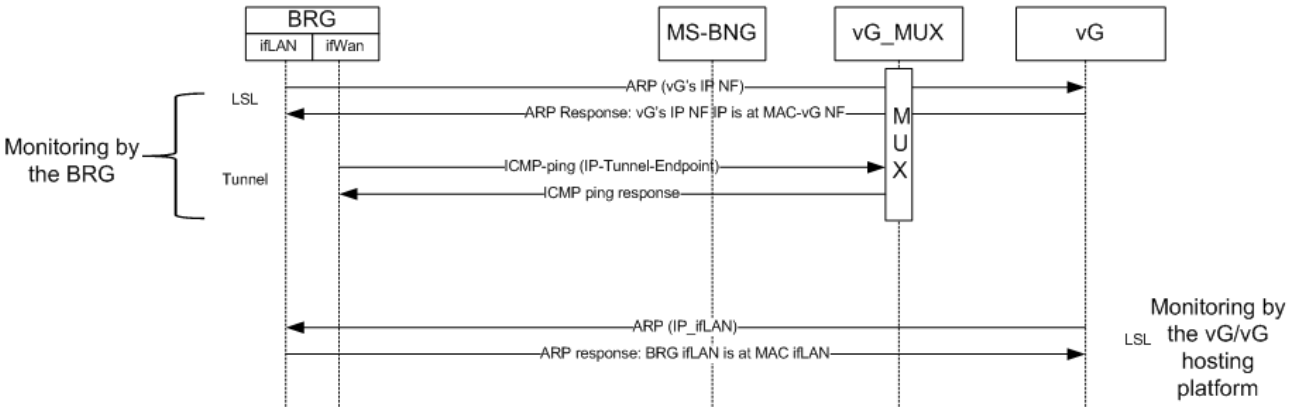

Figure 16 – ARP keep-alive mechanism

7.1.5.1 Connectivity monitoring by the BRG

The BRG monitors the state of its connectivity to the vG, If the vG is not reachable, the BRG checks the state of the LSL, and in the overlay case, the state of the WAN IP session. Consequently:

[R-40] The BRG MUST support the IP session keepalive mechanism described in TR-146 section 6.2.4, using ARP ping for proactive monitoring of the reachability of the default gateway in the vG.

[R-41] The BRG MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861 [12]).
[R-42] The BRG MUST support configuration of the keep-alive time interval, response timeout and the detection multiplier (missed heartbeat count).
[R-43] When the vG is detected as unreachable, the BRG MUST activate its Backup DHCP server and Backup DNS server but continue to test vG reachability.
[R-44] In the overlay LSL case, when the vG is not reachable, the BRG SHOULD be able to query the remote end point of the LSL tunnel (vG_MUX) using ICMP ping.

7.1.5.2 Connectivity monitoring by the vG

The vG monitors the connectivity to its BRG to allow the vG hosting infrastructure to free up some resources when the BRG is not connected.

[R-45] The vG MUST support the IP session keepalive mechanism described in TR-146 section 6.2.4, using ARP ping for proactive monitoring of the reachability of BRG_LAN Interface.
[R-46] The vG MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861)
[R-47] The vG MUST support configuration of the keep-alive time interval, response timeout and the detection multiplier (missed heartbeat count).
[R-48] ARP ping and IPv6 NUD keep-alive mechanisms SHOULD be activated only when no user traffic has been received within a configurable period.
[R-49] The vG_MUX function SHOULD be able to query the BRG WAN interface using ICMP ping.

7.1.6 Traffic classification & Steering

In the case of per device services, selected traffic in the vG may need to use a different group of services/different service chain. This is based on traffic classification and selective steering. As per device services are not in scope of issue 1, this is for further study.

7.1.7 Anti-spoofing

Spoofing where a subscriber attempts to impersonate another subscriber, for example by using the MAC address or the IP address of this other subscriber. In a NERG environment, this would result in a subscriber accessing a vG that is assigned to a different subscriber.

The mapping of a subscriber to a vG is based on the mechanisms described in section 7.1, at two locations, the MS-BNG and the vG_MUX. The role of the MS-BNG differs in the flat and overlay Ethernet models.

In the flat Ethernet LSL model, the MS-BNG maps the subscriber access line to an interface (e.g. MPLS pseudo-wire) toward the vG_MUX. As the subscriber has a dedicated access VLAN, the MS-BNG does not need to look at any MAC address or IP address to send traffic to the vG_MUX.

Similarly, the vG_MUX does not rely on any information inside the traffic issued by the customer, to select the vG so there is no opportunity for spoofing.

In the overlay Ethernet LSL model, the BRG has a WAN IP address, passed by the MS-BNG. At this point, the AAA server has a mapping between the subscriber line identifier (option 82) and the assigned IP address. This unique IP address is used to set up the GRE tunnel toward the vG_MUX. The vG_MUX selects the vG based on the tunnel source IP address. For this reason, it is critical that the broadband network prevents MAC and IP address spoofing. This is covered TR-178, in section 5.4.8 ("Security") for the Access Node and in section 7.1.6 ("Traffic Filtering and QoS") for the MS-BNG. Since the NERG architecture assumes a TR-178 broadband network (access nodes and MS-BNG), there is no additional requirement related to anti-spoofing in this Technical Report.

Note: this Technical Report assumes that the same NSP operates the broadband network, the vG_MUX and the vG. Hence the anti-spoofing done in the broadband network can be trusted.

7.2 Support for Legacy Services

7.2.1 Support for Legacy VoIP service

As explained in section 4.2.3, this section assumes that a VoIP client is hosted by the BRG and uses the IP address assigned to the BRG by the vG's DHCP server. It also assumes that the VoIP traffic is transported over the LSL.

The set of requirements below also assumes the support of PCP by the BRG SIP agent to avoid the need for a SIP ALG on the NAT and drastically reduces the amount of keep-alive messages needed to maintain the NAT binding. The PCP returned lifetime allows appropriate setting of the registration expire timer and therefore avoids offloading both the NAT and the SIP server. The use of PCP also allows successful handling of unidirectional sessions (e.g. access to voicemail or an announcement server).

In the case where the BRG supports a VoIP client and VoIP traffic is transported over the LSL, then the following requirements apply:

[R-50] The BRG MUST support a SIP agent
[R-51] The BRG SIP agent MUST use the IP address assigned to the BRG's LAN interface
[R-52] The BRG MUST forward SIP traffic through the LSL
[R-53] The BRG MUST support PCP (RFC6887 [14]), to instantiate port bindings on the vG's NAT function and to retrieve the assigned public IP address, port number, and lifetime.

The RTP specification (RFC3550 [10]) indicates that RTP uses even ports and RTCP uses odd ports. Also, it specifies that the RTCP port must be set to the RTP port + 1. "a=rtcp" attribute (RFC3605) was specified to relax these constraints but that attribute is not widely implemented. The use of PCP to instruct the NAT function that it must assign a pair of ports that preserve both parity and contiguity will avoid breaking applications without requiring activating an ALG in the NAT or enabling advanced features in the SIP service side (e.g., symmetric RTP/RTCP). The pair

of ports that was retrieved using PCP will be used to build the "c" and "m" lines of an SDP Offer/Answer, hence the following requirements:

[R-54] The BRG MUST support PCP PORT_SET option [I-D-ietf-port-set] to retrieve pairs of ports that preserve the parity and contiguity required for RTP/RTCP sessions. The SIP agent MUST use the returned PCP responses to build its headers (in particular, "Via" and "Contact" fields) and SDP offer/response.

[R-55] The BRG SIP agent MUST disable the "rport" attribute where "rport" is a SIP attribute (RFC3581).

Note: "rport" is used by a SIP user agent (UA) to indicate to a remote SIP UA that it must not use the port number indicated in the topmost Via header to send the response to this request (behavior specified in RFC3261), but it must use the source port number of the request.

[R-56] The BRG SIP agent MUST re-issue a SIP REGISTER message before the expiry of the lifetime returned in the PCP request.

[R-57] The BRG SIP agent MUST support the SDP ALTC attribute [RFC6947] to accommodate dual-stack contexts without requiring any connectivity checks.

7.2.2 Support for Legacy IPTV service

Legacy IPTV service is supported as per TR-101 [2] and places minor additional requirements on the BRG. As explained in section 4.2.3, multicast IPTV traffic is never transported through the LSL.

Figure 17 illustrates the IGMP/MLD model for a BRG using native Ethernet for the LSL uplink (flat model). In this case, there is no WAN IP address on the BRG. The Source IP address of the IGMP requests packets is either the device's IP address or unnumbered (0.0.0.0).

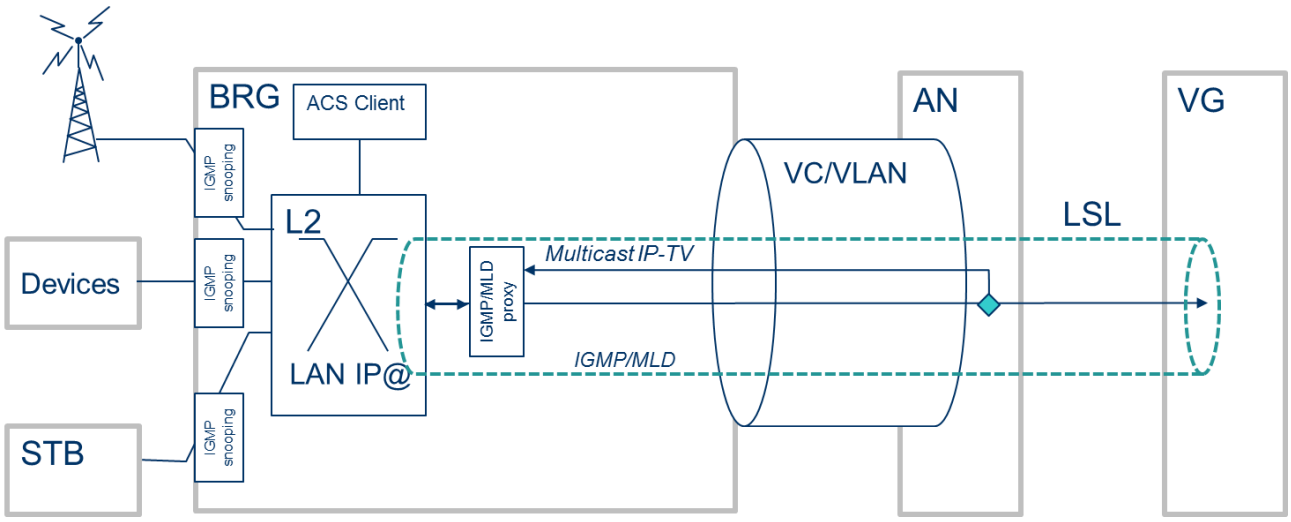

Figure 17 – IGMP/MLD Snooping and Proxy on BRG – Flat LSL

Figure 18 illustrates the IGMP/MLD model for a BRG that implements the LSL as a tunnel (overlay model).

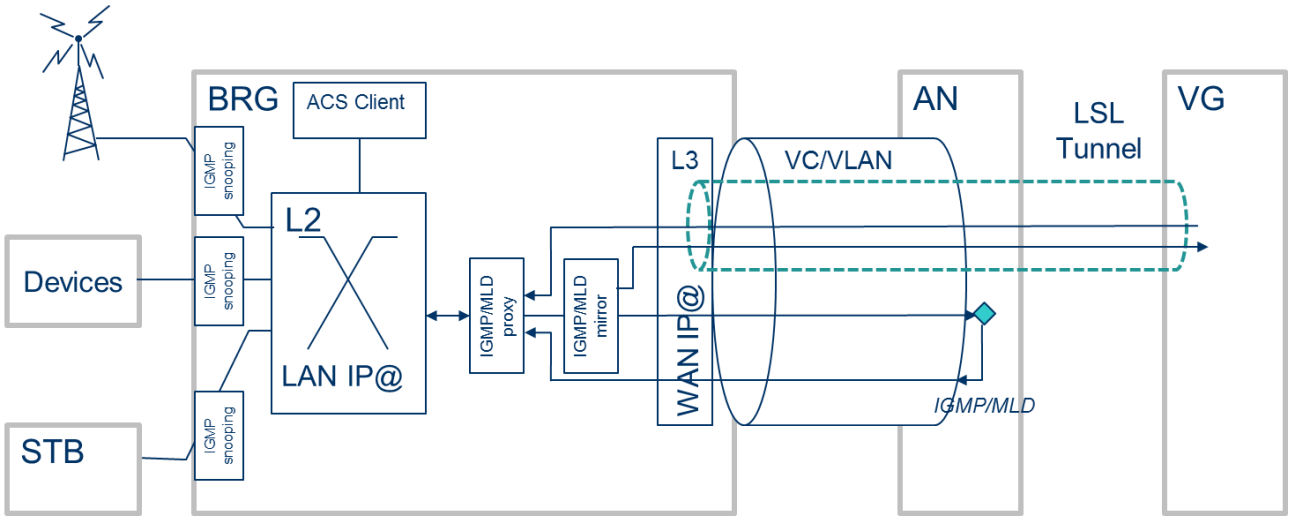

Figure 18 – IGMP/MLD Snooping and Proxy on BRG - Overlay LSL

Port filtering of downstream multicast is implemented on the LAN ports. An IGMP/MLD proxy summarizes the memberships and multicasts it across the set of upstream interfaces.

The consequent requirements on the BRG for the support of an IPTV multicast service are given in section 7.4.3.4.

7.3 QoS

NERG QoS is dependent on the architecture and nodal models. In particular,

- The overlay model creates an additional layer of encapsulation that impacts QoS requirements, as tunneled packets must be prioritized properly.
- When using a BNG as the vG hosting infrastructure, the QoS mechanisms defined in TR-178 [16] are available. Specific requirements are necessary when the vG hosting infrastructure is not co-located with the BNG.

Note: per-device services may trigger the need for an additional level of hierarchical scheduling; however TR-178 [16] does not mandate a particular number of scheduling levels, so it does not create further QoS requirements.

QoS requirements on the BRG for upstream traffic are listed in section 7.4.3.6.
QoS requirements on the vG for downstream traffic are listed in section 7.5.3

7.4 BRG Functional Requirements

The BRG provides functionalities in order to:

- Transfer data between devices within the subscriber's premises

- Transfer data between devices in the subscriber's premises and the network functions that comprise the vG in the NSP network
- Manage the BRG
- Provide local IP address administration when the LSL is unavailable

Section 4.2.1 provides the list of network functions that have been already defined in TR-124 [6] and which are still needed by the BRG. This section defines the additional BRG requirements needed to support the NERG architecture.

7.4.1 General BRG Requirements

[R-58] The BRG MUST comply with the TR-124 GEN.Design requirements for physical design.

[R-59] The BRG MUST comply with the TR-124 REGIONAL requirements for physical design.

[R-60] The BRG MUST comply with the TR-124 GEN.OPS requirements.

[R-61] The BRG MUST comply with the TR-124 GEN.NET requirements for networking protocols.

7.4.2 LAN Requirements

7.4.2.1 LAN Interface Requirements

The BRG provides interfaces to connect to in-home devices. The interfaces on the BRG are chosen to provide the appropriate bandwidth and QoE for the services deployed. For example Internet services are typically provided to connected devices using Ethernet and WiFi. Devices that use Video services may use these same interfaces, or others such as MoCA or HPNA. For all interface types, the BRG needs to be able to:

- Configure the interface
- Discover devices and device topology
- Report the performance of the discovered device
- Execute diagnostics that are specific to the type of LAN interface

TR-124 [6] provides a set of requirements for the functionality provided by LAN interfaces on the BRG described in the IF.LAN section of TR-124 [6] for the different types of LAN interfaces (e.g., HPNA, Ethernet, MoCA).

[R-62] The BRG MUST support a BRG-LANIF functionality that complies with the TR-124 IF.LAN requirements for the respective LAN interfaces.

7.4.2.1.1 WiFi Interface Requirements

The BRG may provide one or more WiFi interfaces. TR-124 [6] provides a set of requirements for the functionality provided by WiFi interfaces and Access Points on the BRG described in the section IF.LAN.WIRELESS.

[R-63] Any WiFi functionality supported by the BRG MUST comply with the TR-124 IF.LAN.WIRELESS requirements.

7.4.2.1.2 Voice ATA Requirements

The BRG may provide a Voice ATA with a corresponding VoIP Client.
TR-124 [6] provides a set of requirements for the functionality provided by Voice ATAs on the BRG described in the section titled IF.LAN.VOICE.ATA. In addition, the BRG may provide additional capabilities necessary for the VoIP Client as described in section 7.2.1. If the BRG provides a Voice ATA, it must comply with the following requirement:

[R-64] The BRG MUST support a Voice ATA that complies with the TR-124IF.LAN.VOICE.ATA requirements.

7.4.2.2 LAN QoS Requirements

The BRG provides the capability to classify, mark and queue traffic that is directed to either the BRG itself or other devices in the subscriber's premises. TR-124 [6] provides a set of requirements for the functionality provided by BRG as described in the section titled LAN.QoS.

[R-65] The BRG MUST support a LAN QoS function that complies with the TR-124 LAN.QoS requirements.

7.4.2.3 BRG's Backup-DHCP Server Requirements

The BRG's DHCP Server (B-DHCP Server) assigns IP addresses to devices within the subscriber's premises when the DHCP Server in the vG is not available as described in section 7.1.3. The BRG's DHCP Server is restricted to serving IP addresses from one IP address subnet because the BRG operates as a bridge device.

TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the sections titled LAN.ADDRESS, LAN.DHCPS and LAN.DHCPv6s.

[R-66] The BRG's B-DHCP Server MUST comply with the following TR-124 LAN.DHCPS requirements: LAN.DHCPS.1,2,4,5,6,7,10,12,13,15,16,17,18.
[R-67] The BRG's B-DHCP Server MUST comply with the TR-124 LAN.DHCPSv6s requirements.
[R-68] The BRG's B-DHCP Server MUST comply with the following TR-124 LAN.ADDRESS requirements: LAN.ADDRESS.1,2,4,5.
[R-69] The BRG's B-DHCP Server MUST comply with the TR-124 LAN.ADDRESSv6 requirements.
[R-70] The BRG's B-DHCP Server's default lease time for DHCPv4 information provided to LAN devices MUST be configurable to a minimum value of no more than 60 seconds. **M**
Note: 10 minutes might be an appropriate typical value.

7.4.2.4 BRG Backup-DNS Requirements

DNS provides the ability to access devices within the home network using a domain name. The BRG's B-DNS provides DNS Server capabilities to devices within the subscriber's premises when the vG DNS Server is not available as described in section 7.1.3.TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the sections titled LAN.DNS and LAN.DNSv6.

[R-71] The BRG MUST support a B-DNS.
[R-72] The BRG B-DNS MUST comply with the following TR-124 LAN.DNS requirements: LAN.DNS.1,2,4,5,6,7,8.
[R-73] The BRG B-DNS MUST comply with the following TR-124 LAN.DNSv6 requirements: LAN.DNSv6.1,2,3,8.

7.4.3 BRG Uplink Requirements

The BRG is a managed device. The data plane traffic is transferred through the LSL interface. The management traffic may or may not be transferred via the LSL interface. Multicast traffic is not transferred over the LSL interface, as explained in section 4.2.3.

7.4.3.1 Uplink Network Interface Requirements

The BRG provides physical interfaces for the BRG to be connected to the Access Network (e.g., xDSL, GPON, Ethernet). The BRG provides the ability to:

- Configure the characteristics of the interface
- Report the performance of the interface
- Execute diagnostics that are specific to the type of uplink network interface

TR-124 [6] provides a set of requirements for the functionality provided by uplink network interfaces on the BRG described in the IF.WAN section of TR-124 [6] for the type of WAN interface.

[R-74] The BRG MUST support a WAN Physical Interface that complies with the TR-124 IF.WAN requirements.

7.4.3.2 Uplink Connectivity Requirements

The BRG's uplink connectivity (i.e., LSL, Management and Control) is considered to be always available during the normal operation of the BRG. TR-124 [6] provides a set of requirements for establishment, tear-down and manual manipulation of uplink connections as described in the section titled WAN.CONNECT as well as the IP interface requirements defined in sections titled WAN.IPv6 and WAN.TRANS.

[R-75] The BRG MUST comply with the TR-124 WAN.CONNECT requirements.
[R-76] The BRG MUST comply with the TR-124 WAN.IPv6 requirements.
[R-77] The BRG MUST comply with the TR-124 WAN.TRANS requirements.

7.4.3.3 Learning Bridge Requirements

The BRG's primary role in the transfer of data is to act a Layer 2 bridge between devices in the subscriber premises and network functions in the vG. TR-124 [6] provides a set of requirements for the functionality provided by a Learning Bridge as described in the section titled BRIDGE.
In addition, the Learning Bridge function is attached to the LSL Interface as described in section 4.2.1 which includes additional requirements for the Learning Bridge.

[R-78] The BRG MUST support a Learning Bridge function that complies with the TR-124 BRIDGE requirements.

7.4.3.4 Multicast Requirements

The BRG supports Multicast functionality that can perform IGMP snooping on LAN interfaces and provides a proxy-routing function that allows Multicast control traffic to be replicated on any tunneled LSL Interface and the network facing interface as described in section 7.2.2.

[R-79] A BRG MUST implement the following IGMP/MLD requirements from TR-124 [6]:
1. LAN.IGMP.BRIDGED Item 1 (snooping)
2. All but item 3 of LAN.IGMP.ROUTED requirements.[1]
3. All but the reference to LAN.IGMP/ROUTED Item 3 in LAN.MLD.ROUTED requirements.

7.4.3.5 OAM Requirements

The BRG provides OAM capabilities as a maintenance end point (MEP). TR-124 [6] provides a set of requirements for the functionality provided by OAM MEPs on the BRG as described in the section titled WAN.ETHOAM.

[R-80] The BRG MUST support an OAM function that complies with the TR-124 WAN.ETHOAM requirements.

7.4.3.6 Upstream QoS Requirements

The BRG provides the capability to classify, mark and queue traffic that is directed to the WAN Interface of the BRG. TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the section titled WAN.QoS.

Note that, although the BRG operates as a bridge, it needs to be able to classify and provide differentiated treatment based on layer 3 fields, as well as any layer 2 classifiers.

[R-81] The BRG MUST support an Upstream QoS functionality that complies with the TR-124 WAN.QoS requirements.

[1] LAN.IGMP.ROUTED item 3 refers to a local firewall that will not exist in a BRG.

The following requirements apply to both the Flat Ethernet and Overlay modes:

[R-82] The BRG MUST support upstream QoS on the WAN port, based on the requirements in section WAN.QoS of TR-124 [6].
[R-83] The BRG MUST support 1 BE queue, 1 EF queue and a minimum of 4 AF queues.

The following requirement only applies to Overlay mode:

[R-84] The BRG MUST support TR-124 issue 4 QoS requirements for Tunneled traffic (section QoS.TUNNEL).

7.4.3.7 DHCP Client Requirements

The BRG provides the capability to obtain its IP address(es) used for the device (BRG as a host) and its WAN interface. In addition, the BRG can obtain other information through the use of DHCP options in order to configure the LSL Interface as defined in section 7.1 (overlay LSL case) and to configure the Management Client as defined in section 7.4.4.
TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the section titled WAN.DHCPC.

[R-85] The BRG MUST support a DHCP Client that complies with the TR-124 WAN.DHCPC requirements.
[R-86] The BRG's DHCP client MUST insert the Device-Type information in the dedicated Option 125 sub-option 24 in its all its DHCP messages
[R-87] The default value of DHCP option Device-Type MUST be the string "BRG" (without the quotes)
[R-88] The value of DHCP option Device-Type MUST be configurable by the NSP.

7.4.3.8 PCP Client Requirements

The NAT function is located on the vG. A PCP client embedded in the BRG allows the opening of pinholes to avoid the need for ALGs.

[R-89] The BRG MUST support a PCP Client based on the Port Control Protocol [RFC 6887] so as to instantiate port forwarding rules on the vG NAT and to retrieve the assigned public IP address, port number and lifetime.

7.4.3.9 Authentication Client Requirements

The BRG provides the capability for itto be authenticated within the operator network.
TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the section titled WAN.dot1x. Additional authentication mechanisms can be found in TR-146.

[R-90] The BRG MUST comply with the TR-124 WAN.dot1x requirements.
[R-91] The BRG SHOULD comply with the authentication requirements in TR-146.

7.4.4 BRG Management Client Requirements

The BRG's Management Client is a TR-069 [1] client that provides the capabilities to configure and monitor the BRG. TR-124 [6] provides a set of requirements for the functionality provided by the BRG as described in the sections titled MGMT.GEN and MGMT.REMOTE.TR-069. The Management system is either accessed through the BRG's WAN IP interface or the BRG-LSL interface.
The requirements for providing access to the Management Client through the BRG-LSL interface are defined below:

[R-92] The BRG MUST support a Management Client that complies with the TR-124 MGMT.REMOTE.TR-069 requirements.
[R-93] The BRG MUST be able to be TR-069 managed through the LSL, thus through NAT, using PCP.
[R-94] The BRG MUST comply with the TR-124 MGMT.GEN requirements.

7.4.5 Network Time Protocol Requirements

Several functions (e.g. the Management Agent) within the BRG need to use the device's time capabilities. The BRG obtains the time using the Network Time Protocol (NTP). TR-124 [6] provides a set of requirements for this functionality in the section titled MGMT.NTP.

[R-95] The BRG MUST support a NTP Client that complies with the TR-124 MGMT.NTP requirements.

7.4.6 User Interface Requirements

The BRG's User Interface provides the ability to manage the functions hosted on the BRG via a Graphical User Interface (GUI). Access to the GUI is not provided across the BRG's WAN interface; only local access is provided.
Access to the User Interface requires the user to be authenticated, e.g.: a login/password.

[R-96] The BRG MUST support a Graphical User Interface.
[R-97] The User Interface MUST authenticate the user.

7.5 vG Functional Requirements

7.5.1 vG LAN DHCP Requirements

7.5.1.1 IP addressing

In the NERG architecture, L3 functions including IP addressing mechanisms and routing/NAT are part of the vG. This section provides the requirements on IP assignment for both IPv4 and IPv6.

[R-98] The vG MUST support IPv4/IPv6 dual stack functionality. **M**

7.5.1.1.1 IPv4 addressing

The vG is responsible for IPv4 address assignment. The LANs behind different subscribers' BRGs may use overlapping private IP addresses. Each LAN device will obtain a unique private address from the subnet. The vG DHCP server may also assign IP addresses to hosts within the extended LAN (vNAS, M2M, etc.). The private addresses are then translated via the vG NAT function to provide internet access. Therefore, the following requirements need to be met:

[R-99] The vG MUST support a DHCP server that complies with the TR-124 LAN.DHCPS (1,2 4, 6, 7, 8, 9, 11, 12, 13, 14, 15, 16, 17, 18) requirements wherein the word "Device" must be read as "vG".

[R-100] The vG MUST provide application layer support for host name mapping, booting, and management including DHCPv4 and the Domain Name System (DNS) protocol. This includes support for the standards below and required elements of associated updates:

- RFC 1034 Domain Names – Concepts and Facilities
- RFC 1035 Domain Names – Implementation and Specification
- RFC 2131 Dynamic Host Configuration Protocol
- RFC 2132 DHCP Options and BOOTP Vendor Extensions
- RFC 2181 Clarifications to the DNS Specification
- RFC 2939 Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types

[R-101] The vG MUST act as a DHCPv4 server to local LAN devices, supporting all LAN devices.

[R-102] The vG MUST support a minimum of 253 LAN devices[2].

[R-103] The vG MUST support turning off the DHCPv4 server via a configuration change and a user interface (see 7.5.12).

[R-104] The DHCPv4 server functionality of the vG MUST verify that an address is not in use prior to making it available in a lease (e.g. via ping or ARP table validation) even when lease information shows that it is not in use.

[R-105] The default lease time for DHCPv4 information provided to LAN devices MUST be configurable. The default value MUST be 24 hours, or an operator-specific value

[R-106] When the domain name that the DHCPv4 server passes to LAN devices has not been set, the value "domain_not_set.invalid" SHOULD be used.

[R-107] The vG SHOULD allow the DHCPv4 server to be configured so that specific MAC addresses can be identified as being served or not served.

[R-108] The vG SHOULD allow the DHCPv4 server to be configured with a default setting (provide IPv4 addresses or not provide IPv4 addresses) for devices whose MAC addresses have not been specified in accordance with [R-107].

[R-109] The DHCPv4 server functionality of the vG SHOULD provide a mechanism by which an IPv4 address can be assigned to a particular LAN device by MAC address. The user interface to establish this association may use an alternate mechanism to identify this assignment (e.g. by selecting the device using its current IPv4 address or device name) and the MAC address may be transparent to the user. These addresses may include addresses

[2] In average, the number of devices in on customer's LAN is expected to be much lower than 253.

within the default subnet or addresses from additional public/private subnets that may be provisioned.

[R-110] When each vG has its own DHCPv4 server, it MUST be able to assign a private IPv4 address to each IP host. **M**

[R-111] When there is a common DHCPv4 server for a set of vGs, it MUST allow using the same private IPv4 subnet across all subscribers, to assign a private IP address to each IP host.

[R-112] The vG's DHCP Server SHOULD be able to add other subnets to a BRG.

[R-113] The vG's DHCP Server MUST be able to assign public IPv4 addresses to selected IP hosts.

7.5.1.1.2 IPv6 addressing

To support DHCPv6, the vG must satisfy the following requirements:

[R-114] The vG MUST support DHCPv6 server messages and behavior as per RFC 3315 with required elements from RFC updates.

[R-115] The vG MUST support and be configurable to enable/disable address assignment using DHCPv6.

[R-116] The vG MUST have an algorithm and/or allow configuration as to which /64 prefix to use, from any assigned prefixes or its own ULA prefix.

[R-117] The vG SHOULD be configurable to support rules as to which host devices will be assigned addresses through DHCPv6. It should be possible for a service provider to place their own host devices at the customer premises and have the vG only support DHCPv6 address assignment to those devices. Note that this does not require use of the RA "M" flag, as the service provider host devices can be configured to always use DHCPv6 for address assignment. The DUID may help to identify host devices.

[R-118] The vG MUST be configurable to enable/disable prefix delegation via DHCPv6.

[R-119] The vG MUST support delegation of any received WAN prefix and its own ULA prefix, that is shorter than /64, using the mechanisms of RFC 3633 with required elements from RFC updates.

[R-120] The WAN / ULA prefixes that a vG is allowed to further delegate SHOULD be configurable.

[R-121] The vG MUST support DHCPv6 Information request messages.

[R-122] The vG MUST support the following DHCPv6 options: IA_NA (RFC 3315), IA_PD (RFC 3633), and DNS_SERVERS (RFC 3646).

[R-123] The vG SHOULD support Reconfigure Accept (RFC 3315) and pass the additional set of DHCP options received from the DHCP client on its WAN interface to IPv6 hosts.

[R-124] The options that the vG provides via DHCPv6 MUST be configurable.

[R-125] If address selection policy option is requested in a DHCPv6 request from hosts, the vG SHOULD advertise the generated address selection policy.

[R-126] The vG MUST be able to assign a unique, persistent /128 IA_NA address to individual IPv6 hosts. **M**

[R-127] The vG MUST be able to assign addresses from a single IPv6 prefix to a set of devices connected to the same BRG and the hosts connected to the same Extended LAN. **M**

[R-128] The vG MUST support multiple prefixes to subsets of devices connected to the same BRG and to the hosts connected to the same Extended LAN. **M**

In the case of SLAAC, the vG must satisfy the following 2 requirements:

[R-129] The vG MUST be able to advertise a single prefix for use among a set of devices connected to the same BRG and the hosts connected to the same Extended LAN. **M**

Unlike DHCP subscribers, SLAAC subscribers are not stateful and it is difficult to track the connectivity of the device, hence:

[R-130] The vG MUST support Neighbor Unreachability Detection (NUD), as defined in RFC 4861, "Neighbor Discovery in IPv6", to verify if a device is still connected to the network.

7.5.2 vG DNS Requirements

DNS provides the ability to access devices within the home network using a domain name (e.g.: user-friendly printer names, user interface access).

7.5.2.1 IPv4 DNS requirements

[R-131] The vG MUST be capable of acting as a DNS server to LAN devices, passing its address as the DNS server back to these devices in DHCPv4 requests.

[R-132] The vG SHOULD allow the user to specify that either network-learned or user-specified addresses be passed back to LAN devices as the DNS server(s) in DHCPv4 responses, instead of the vG's address.

[R-133] The vG MUST be capable of acting as a DNS server to LAN devices, passing its address as the DNS server back to these devices in DHCPv4 requests.

[R-134] When the vG learns DNS name server addresses from multiple WAN connections, the vG MUST follow specified DNS selection policy (if one is configured) to make recursive queries to DNS name servers, or (if there is no DNS selection policy) MUST query a server on each connection simultaneously and provide the requesting LAN client with the first returned positive result from these DNS servers. A negative response MUST not be transmitted to a LAN device until all WAN DNS servers have either timed out or returned a negative response to a common query.

Service providers may choose not to provide DNS name server addresses on certain connections in a multiple connection configuration.

[R-135] The vG SHOULD support additional DNS entries, as there could be additional types of CPE.

[R-136] The vG MUST be able to maintain local DNS entries for a minimum of 253 local LAN devices. This information can be obtained through auto discovery (e.g. from DHCPv4 requests, such as Client Identifier, and other protocol information). When unknown, the entry MUST be of the form "unknownxxxxxxxxxxxx" where "x" represents the MAC address of the associated LAN device.

[R-137] The vG SHOULD provide a configuration and user interface mechanisms to override the learned names of all LAN devices except that of the vG itself.

[R-138] If the vG's DNS server is implemented as a forwarding proxy, it MUST be done according to the recommendations in RFC 5625.

7.5.2.2 IPv6 DNS Requirements

[R-139] The vG MUST act as a DNS server for IPv6-capable LAN devices by supporting IPv6 (AAAA) records in its DNS server (as per RFC 3596) and allowing these records to be queried using either IPv4 or IPv6 transport (per RFC 3901).

[R-140] The vG MUST attach all known (for the host device) globally scoped IPv6 addresses to the DNS record for a particular host device (see TR-124 LAN.DNS.6), as AAAA records for that device.

[R-141] The vG SHOULD support dynamic DNS (DDNS) for devices to provide their own DNS information. This would override any DNS entries the vG might have created for the IP addresses included in the DDNS request.

[R-142] The vG MUST be able to query for A and AAAA records using either IPv4 or IPv6 transport to DNS recursive name servers in the WAN.

[R-143] The vG SHOULD use a DNS recursive name server obtained through DHCPv6 option 23 (OPTION_DNS_SERVERS) to query for AAAA records to the WAN, as its first choice.

[R-144] When the vG is proxying DNS queries for LAN devices, it SHOULD use IPv6 transport regardless of the transport mode used by the LAN device, when querying to the WAN. This is only possible if the vG has IPv6 addresses for DNS recursive name servers on the WAN.

[R-145] The vG MUST support receiving at least 2 DNS recursive name server IPv6 addresses from the network through DHCPv6 option 23 (OPTION_DNS_SERVERS) (RFC 3646).

[R-146] The vG SHOULD allow the user to specify that the network-learned or user-specified DNS recursive name server addresses be passed back to the LAN devices in DHCPv6 responses instead of the vG's address itself as the DNS recursive name server(s).

[R-147] When the vG learns DNS name server addresses from multiple WAN connections, the vG SHOULD make a recursive query to the DNS name server specified with DNS selection policy that is obtained through DHCPv6 (RFC 6731 [13]) or a manually configured DNS selection policy.

7.5.3 Downstream QoS

A traditional RG aggregates multiple end user devices and provides a single IP session to the BNG. The MS-BNG will then apply QoS to that single IP session. For a vG this is different, as the end devices are directly connected. Therefore QoS is no longer applied to a single subscriber aggregate, but a group of hosts associated with a single subscriber. This fundamental difference requires new QoS functionality in the vG.

In the scenario where the vG is implemented within the MS-BNG, the MS-BNG can perform hierarchical scheduling as per the requirements in TR-178 [16] (sections 4.4.3 [Options: Hierarchical QoS] and 7.1.4 [Hierarchical QoS on MS-BNG]).

In the scenario where the vG is deployed remotely from the MS-BNG, the vG needs to include a traffic management function (TMF) that both rate limits the aggregate of downstream traffic from the vG, and provides differential queuing. The rate limit value used is a configured rate but may be dynamically modified via IGMP joins/leaves to adjust for any AN inserted multicast streams, as per TR-178 [16].

7.5.3.1 Classification Requirements

[R-148] The vG MUST be able to classify traffic based on the following fields: src/dst IP, src/dst port, DSCP/ToS, protocol number and Ethernet p-bits. **M**
[R-149] The vG MUST support traffic classification based on layer 2-4 information. **M**
[R-150] The vG MUST support DiffServ and Ethernet priority marking/remarking on the basis of classification. **M**

The following requirements apply only to the Overlay mode:
[R-151] The vG_MUX function MUST copy DiffServ and (when appropriate) 802.1p markings to the tunnel header. **M**

7.5.3.2 Queuing and Policing requirements

[R-152] The vG_MUX MUST support queuing and policing per IP stack (i.e. separately for IPv4 and IPv6) and per IP host (across both IPv4 and IPv6). **M**
[R-153] The vG_MUX MUST support downstream and upstream rate limiting of the LSL. **M**

7.5.3.3 Scheduling Requirements

[R-154] The vG and vG_MUX hierarchical scheduler MUST support the following types of scheduling: strict priority, weighted round robin, and round robin. **M**
[R-155] The vG_MUX MUST be able to aggregate all queues and policers to an aggregate rate in both ingress and egress directions. **M**
[R-156] The vG_MUX MUST be able to assign a minimum bandwidth guarantee to a queue. Note: the minimum can be zero (no guarantee). **M**
[R-157] The vG_MUX MUST be able to map traffic classes to queues. **M**
[R-158] Queues MUST be able to be configured with a given forwarding behavior such as Expedited Forwarding, Assured Forwarding, etc. **M**
[R-159] The vG MUST be able to inspect user IGMP requests in order to track multicast usage and to monitor bandwidth on the AN. **M**
[R-160] The vG MUST be able to modify rate-limit values using configured information that associates bandwidth with IGMP group membership. **M**

7.5.4 Device Inventory

An end-user device may be identified based on a combination of a permanent (e.g. subscriber line-identifier + MAC address) and a user friendly identification mechanism, which adds attributes to a device. Examples of such attributes include:

- Device name (e.g. “living-room TV”)
- Device owner (e.g., “Little John”)
- Device type (e.g. “Tablet”)

Devices that share a common attribute value are said to be part of the same class of devices (e.g., all TVs in the home network).
The tagging mechanism may be:

- Automated (e.g. device fingerprinting in the network, based on MAC address, vendor-id, DHCP attributes, UPnP, etc.)
- Manual (subscriber inputs device attributes via a self-care portal)
- Hybrid (e.g., subscriber complements attributes based on suggestions derived from fingerprinting)

The device identity information is expected to be stored in a repository. This information needs to be accessible by the subscriber, through a graphical user interface.

Note: How device identification is retrieved, formatted and stored, is for further study.

[R-161] The vG MUST support creating an inventory of connected devices in the home network, combining user administered device information (e.g. device name) and dynamic network information (e.g. IP address).
[R-162] The vG device inventory MUST be accessible to both the subscriber and the network service provider.
[R-163] The vG addressing mechanisms MUST support device specific address assignment, where a fixed IP address or a specific pool of IP addresses are used for a given device or class of devices. **M**
[R-164] The vG MUST support default behaviors for unidentified devices (e.g., accept, accept with restrictions, force identification, deny). **M**
[R-165] The vG MUST support enabling remote (i.e. out-of-home) access to specific devices (see section 7.5.12 for the underlying mechanisms).
[R-166] The vG MUST support a configurable limit on the number of simultaneously connected devices, in order to protect the vG host against attacks or device malfunctions. **M**

7.5.5 NA(P)T Requirements

The vG performs NA(P)T in order to provide connectivity between LAN devices and external resources (e.g. other subscribers, the Internet, external applications). The vG NAT function must support flexible control of port forwarding rules, in order to allow external users or applications to access end-user devices. Configuration may be static (the subscriber enabling remote access for a specific end-user device by provisioning port forwarding on the self-care portal / vG UI) or dynamic (application driven).

[R-167] The vG MUST support Port Forwarding as part of the NAT function that complies with the TR-124 LAN.PFWD (1,2,3,4) requirements wherein the word “Device” must be read as “vG”. **M**
[R-168] The NAT/PAT function of the vG MUST support port-forwarding rule configuration using PCP. **M**
[R-169] The NAT/PAT function of the vG MUST support port-forwarding rule configuration using NAT-PMP. **M**
[R-170] The NAT/PAT function of the vG MUST support port-forwarding rule configuration using UPnP-IGD. **M**
[R-171] The NAT/PAT function MUST be able to differentiate subscribers with overlapping private IPv4 addresses. **M**

[R-172] The vG MUST provide an IPv4 NAT function in one of the following ways:
a. all IP hosts belonging to one vG are NATed to a single, unique public IPv4 address
b. all IP hosts belonging to multiple vGs are NATed to a single public IPv4 address.

7.5.6 DDoS Prevention Requirements

The vG provides DDOS protection to the NERG functions and end-user devices for known DDoS attacks that are initiated from the vG's WAN interface side.

[R-173] The vG DDoS function MUST provide DDoS protection for the NERG and end-user devices including protection from ping of death, SYN flood, LAND and variant attacks.
[R-174] The vG DDoS function MUST reject packets from the WAN with source MAC addresses of hosts in the Extended LAN or end-user devices.
[R-175] The vG DDoS function MUST reject packets from the WAN with invalid IP (v4 or v6) addresses (e.g. broadcast addresses or IP (v4 or v6) addresses matching those assigned to hosts in the Extended LAN or the subscriber's home network.
[R-176] The vG DDoS function MUST reject any unidentified Ethernet packets (i.e. any packet that is not associated with IP (v4 or v6).
[R-177] The vG DDoS function MUST perform anti-spoofing filtering for IPv6. All IPv6 traffic sent to the WAN from the LAN MUST have an IPv6 source address with a prefix assigned to the home network.
[R-178] When the vG DDoS function performs anti-spoofing filtering for IPv6, the vG DDoS function MUST filter all upstream IPv6 traffic from the home network until the IPv6 prefix is assigned to the home network.

7.5.7 ALG Requirements

The vG provides ALGs to hosts in the Extended LAN and end-user devices in order to facilitate communication with other entities across the vG's WAN interface. While various functions utilize PCP to alleviate the need for ALGs, other functions and end-user devices still require assistance to communicate with entities across the vG's WAN interface.

[R-179] The vG ALG MUST allow pass-through of IPv4 traffic in which the payload is compressed or encrypted (e.g. VPN traffic) that originates or is directed to hosts in the Extended LAN or end-user devices.
[R-180] The vG ALG MUST allow hosts in the Extended LAN and end-user devices to initiate IPv4 IPsec sessions to an external network. This function MUST also work through the vG's NAPT function.
[R-181] The vG ALG MUST support UDP encapsulation of IPv4 IPsec packets from hosts in the Extended LAN and end-user devices as defined in IETF RFC 3948.
[R-182] The vG ALG MUST support negotiation of NAT traversal with IKE as identified in IETF RFC 3947 from hosts in the Extended LAN and end-user devices.
[R-183] The vG ALG MUST allow multiple hosts in the Extended LAN or end-user devices to launch independent and simultaneous IPv4 IPsec sessions. These sessions can be to the same or different destinations.

[R-184] The vG ALG SHOULD support a minimum of 4 concurrent IPv4 IPsec sessions per host in the Extended LAN or end-user device. These sessions can be to the same or different destinations.

[R-185] The vG ALG MUST seamlessly handle RTSP traffic to hosts in the Extended LAN and end-user devices with no user intervention.

[R-186] The vG ALG MUST allow the service provider to administratively enable and disable ALG functionality (e.g. disable the SIP ALG) on a per host in the Extended LAN or end-user device basis.

[R-187] If the vG does not maintain SIP clients using PCP, the vG ALG MUST be aware of the presence of active SIP clients on hosts in the Extended LAN and end-user devices using rules (e.g. matching IP address, port, or protocol number through interception of SIP REGISTER messages).

[R-188] If the vG does not maintain SIP clients using PCP, the SIP ALG function MUST keep track of SIP events (e.g. REGISTER reply from the registrar) and maintain allocated resources within the event timeout period.

[R-189] The vG ALG MUST allow hosts in the Extended LAN and end-user devices to originate IPv4 FTP sessions to an external network. This function MUST also work through the vG's NAPT function.

7.5.8 Firewall Requirements

The vG provides firewall functionality to the hosts in the Extended LAN and end-user-devices in order to protect incoming communication from entities through the vG's WAN interface.

[R-190] The vG Firewall MUST drop or deny IPv4 access requests from WAN side connections to hosts in the Extended LAN and end-user devices, except in direct response to outgoing traffic or as explicitly permitted through configuration via the vG UI (e.g. for port forwarding or management).

[R-191] The vG Firewall MUST support a separate firewall log to maintain timestamped records of transactions according to firewall rules.

[R-192] The vG Firewall SHOULD provide protection against the following:

- Port scans
- Packets with same source and destination addresses
- Packets with a broadcast source address
- Downstream packets with a LAN source address
- Invalid fragmented IP (v4 or v6) packets
- Fragmented TCP packets
- Packets with invalid TCP flag settings (NULL, FIN, Xmas, etc.)
- Fragmented packet headers (TCP, UDP and ICMP)
- Inconsistent packet header lengths
- Packet flooding
- Excessive number of sessions
- Invalid ICMP requests
- Irregular sequence differences between TCP packets

[R-193] Each type of attack for which protection is provided SHOULD be configurable on the vG Firewall and be on by default.

[R-194] The vG Firewall MUST support passing and blocking of traffic by user-defined and service provider configurable rules.
[R-195] The vG Firewall MUST support setting firewall rules by the service provider that cannot be altered or overridden by the user.
[R-196] The vG Firewall MUST support the user temporarily disabling specific user-defined rules or all user defined rules, that is, without deleting the rules.
[R-197] The vG Firewall MUST support the user specifying the order in which firewall rules are processed.
[R-198] The vG Firewall SHOULD support specification of any of the following in a firewall rule:

- destination IP (v4 or v6) address(es) with subnet mask
- originating IP (v4 or v6) address(es) with subnet mask
- source MAC address
- destination MAC address
- protocol (0-255, or by alias: TCP, UDP, ICMP, IP, IGMP, gre, ipinip, pim, nos, ospf, …)
- source port
- destination port
- IEEE 802.1Q user priority
- FQDN (fully qualified domain name) of WAN session
- DiffServ codepoint (IETF RFC 3260)
- Ethertype (IEEE 802.3) length/type field)
- Traffic matching an ALG filter
- IEEE 802.1Q VLAN identification
- packet length
- TCP flags (urg, ack, psh, rst, syn, fin)
- IP option values (potentially name aliases)
- logical interface of source
- logical interface of destination

[R-199] The vG Firewall MUST allow for configuration of basic firewall profiles (e.g., Off, High, Medium, and Low).
[R-200] The vG firewall SHOULD be either ICSA certified (*www.icsalabs.com*) or be able to display all the attributes necessary for ICSA certification for the current version of either the Residential category or the Small/Medium Business (SMB) category.

7.5.9 Support for Value Added Services

Support for these services, such as M2M services, (shared) DLNA services etc.… is for further study, and is out of scope of this document.

7.5.10 Ethernet Service Extension – Extended LAN

Ethernet service extension essentially extends the connectivity of the user's home LAN to the network in order to provide access to additional services. The broadcast domain is extended to allow Ethernet communication between these distributed LAN hosts and devices.

A logical Ethernet switch within the vG connects the subscriber home network and the NERG functions (NFs) hosted in the vG; an example of how this can be done is shown on Figure 19. These NFs can host applications to provide value added services such as the ones listed in Section 5. The logical Ethernet Switch provides a private LAN to each vG.

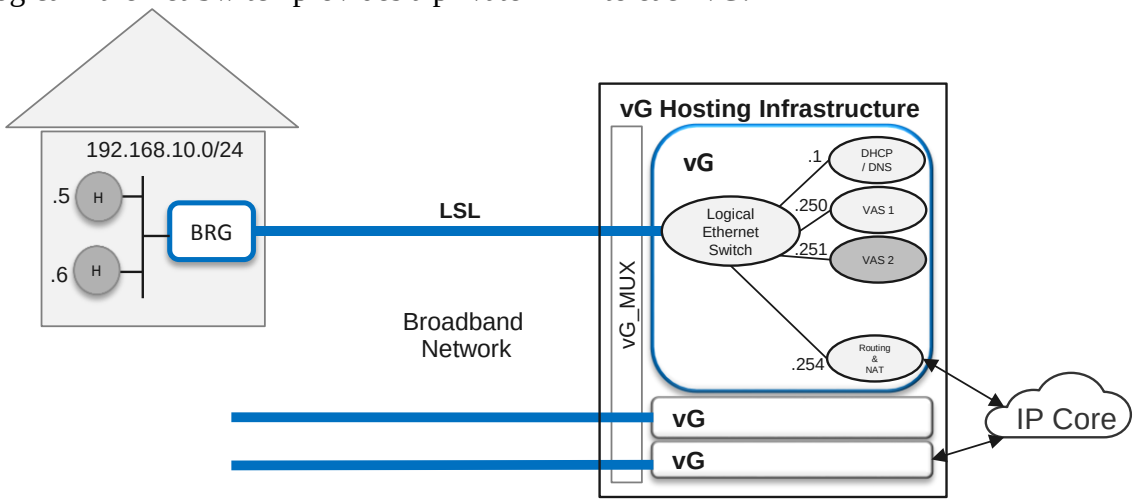

Figure 19 – LAN Extension

[R-201] The vG MUST support a logical Ethernet switch connecting the in-home hosts (reachable over the vG-LSL-Interface), the physical/virtual hosts in the vG and the routing/NAT (default gateway).

[R-202] The vG logical Ethernet switch MUST comply with IEEE 802.1Q specifications with respect to broadcast, multicast and MAC learning.

7.5.11 Management Client Requirements

The vG Management Client is a mandatory function that provides fault, configuration and performance monitoring capabilities and management of application layer functions. A vG can have multiple vG Management Clients where a vG Management Client manages one or more functions hosted within the Service Provider's network.

A vG Management Client instance can manage functionality for one or more vGs.

A vG Management Client instance can be hosted within a vG Hosting infrastructure (e.g., hosted by the MS-BNG to manage the NAT functions) or a vG Management Client instance can be hosted as a function within the vG.

The protocol and information elements (e.g., NETCONF/Yang, TR-069) used by the vG Management Client are beyond the scope of this document.

[R-203] For each function within the vG, the NERG MUST support a vG Management Client to manage the fault, configuration, performance monitoring capabilities and management of the function's application layer.

7.5.12 User Interface Requirements

The vG-UI is a mandatory functionality that allows management of the functions hosted within the Service Provider's network, as well as on the BRG. Access to perform operations on these functions is provided to the following (user) roles:

- One (or more) end-users with admin rights
- Other end-users
- Service Provider Technician
- Service Provider Help Desk

The vG-UI is a Graphical User Interface (GUI) provided as a Web-based portal.

Access to the vG-UI requires the user to authenticate. When performing an operation on a function, the vG-UI ensures that the user/role is authorized to perform that operation.

The vG-UI can be deployed as a shared Web portal between Subscribers or as a Web portal deployed for a specific Subscriber. In either case, the vG-UI can be virtualized.

[R-204] The NERG MUST support a vG-UI as a Web-based portal.
[R-205] The vG-UI MUST authenticate the vG-UI user and role prior to allowing access to the vG-UI. **M**
[R-206] The vG-UI MUST only allow authorized actions for a given user/role.
[R-207] The vG-UI MUST support access to the admin user and the end-users through the BRG.
[R-208] The vG-UI MUST support access to the Service Provider's technicians through the Service Provider network.
[R-209] The vG-UI MUST support access to the Service Provider's help desk through the Service Provider network.

End of Broadband Forum Technical Report TR-317