TECHNICAL REPORT

TR-328
Virtual Business Gateway

Issue: 1
Issue Date: July 2017

© The Broadband Forum. All rights reserved.

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

 © The Broadband Forum. All rights reserved

RIGHTS ARE DESCRIBED ON THE TECHNICAL REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

The text of this notice must be included in all copies of this Technical Report.

© The Broadband Forum. All rights reserved

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	3 July 2017	21 July 2017	Ron Insler, RAD Guiu Fabregas, Nokia	Original

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

Editor	Ron Insler	RAD	ron_i@rad.com
	Guiu Fabregas	Nokia	guiu.fabregas@nokia.com
SDN & NFV Work Area Director(s)	George Dobrowski	Huawei Technologies	georgedobrowski@mail01.huawei.com
	Ken Ko	Adtran	ken.ko@adtran.com

© The Broadband Forum. All rights reserved

TABLE OF CONTENTS

 © The Broadband Forum. All rights reserved

 © The Broadband Forum. All rights reserved

List of Figures

List of Tables

 © The Broadband Forum. All rights reserved

 © The Broadband Forum. All rights reserved

Executive Summary

TR-328 specifies the Virtual Business Gateway (VBG) System architecture. The VBG System consists in virtualizing some of the functionalities of a Business Gateway (BG) to a flexible hosting environment which may be located at the customer premises, in the operator's network or using a combination of the two.

With the VBG System architecture, the functions provided traditionally by the BG are now distributed between a simplified on-site physical device called the pBG (physical Business Gateway) and a virtualized component, called vBG (virtual Business Gateway). The vBG hosting environment can benefit both from network equipment and recent network virtualization technology.

The Technical Report describes the motivations to deploy the VBG System architecture, based on the use cases that it enables. In particular, it facilitates simplification of the customer located equipment, customer self-provision through a portal, rapid introduction of new services, decommissioning of unsuccessful ones, and upselling value-added services. All without the need to deploy specialized hardware devices to remote enterprise sites. Examples of value-added services include: enterprise class firewall, Wide Area Network (WAN) optimization, etc.

Following a high level architecture description and some examples of deployment models, TR-328 defines the following set of technical requirements:

- End-to-end network requirements and support for existing as well as new business services
- pBG requirements
- vBG requirements

The target audience for this document is:

- Service Providers (SPs) who want to virtualize the Business Gateway,
- Suppliers who want to build interoperable pBG’s and vBG’s,
- System integrators responsible for the integration of VBG System services in the SP’s information systems.

 © The Broadband Forum. All rights reserved

1 Purpose and Scope

1.1 Purpose

This Technical Report specifies architecture and requirements for the virtual business gateway. The virtual business gateway architecture describes the migration of functionalities running on a business gateway to the network service provider's infrastructure for enabling network-based features and services. Such migration is expected to simplify the deployment and management of network and business services.

TR-328 targets different business premises sizes, e.g., small and medium enterprises (SMEs), campus, as well as single office and home offices (SOHO).

1.2 Scope

The scope of the Technical Report includes:

- Use cases and business drivers from both customer and operator sides
- BG functional decomposition and impact on existing reference points
- Defining a set of network architectures in support of the VBG System
- Defining new reference points, where applicable
- Enabling multi-homed sites (e.g., backup)
- Defining requirements for performance monitoring
- Identifying high level management and orchestration needs
- Analyzing impact on security and privacy

The following aspects are not in scope of the current issue of TR-328:

- SDN control
- Detailed policy aspects and attributes
- Detailed management requirements and portal capabilities
- VBG System data models
- vBG redundancy (1:1, N:1, etc.)
- Two or more redundant pBG's in a business branch

 © The Broadband Forum. All rights reserved

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [17].

MUST	This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course.
SHOULD NOT	This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label.
MAY	This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

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]	TR-124 Issue 5	*Functional Requirements for Broadband Residential Gateway Devices*	BBF	2016

© The Broadband Forum. All rights reserved

[4]	TR-134 Cor. 1	*Broadband Policy Control Framework (BPCF)*	BBF	2013
[5]	TR-146	*Subscriber Sessions*	BBF	2013
[6]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[7]	TR-181 Issue 2 Amd. 11	*Device Data Model for TR-069*	BBF	2016
[8]	TR-317	*Network Enhanced Residential Gateway*	BBF	2016
[9]	TR-359	*A Framework for Virtualization*	BBF	2016
[10]	TR-390	*Performance Measurement from IP Edge to Customer Equipment using TWAMP Light*	BBF	2017
[11]	802.1Q	*Bridges and Bridged Networks*	IEEE	2014
[12]	802.1X	*Port Based Network Access Control*	IEEE	2004
[13]	802.3	*CSMA/CD access method and physical layer specifications*	IEEE	2005
[14]	RFC 951	*Bootstrap Protocol*	IETF	1985
[15]	RFC 1195	*Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*	IETF	2007
[16]	RFC 1918	*Address Allocation for Private Internets*	IETF	1996
[17]	RFC 2119	*Key words for use in RFCs to Indicate Requirement Levels*	IETF	1997
[18]	RFC 2131	*Dynamic Host Configuration Protocol*	IETF	1997
[19]	RFC 2132	*DHCP Options and BOOTP Vendor Extensions*	IETF	1997
[20]	RFC 2328	*OSPF Version 2*	IETF	1998
[21]	RFC 2474	*Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*	IETF	1998
[22]	RFC 2475	*An Architecture for Differentiated Services*	IETF	1998
[23]	RFC 2597	*Assured Forwarding PHB Group*	IETF	1999
[24]	RFC 2784	*Generic Routing Encapsulation (GRE)*	IETF	2000
[25]	RFC 2939	*Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types*	IETF	2000
[26]	RFC 3046	*DHCP Relay Agent Information Option*	IETF	2001
[27]	RFC 3246	*An Expedited Forwarding PHB (Per-Hop Behavior)*	IETF	2002
[28]	RFC 3260	*New Terminology and Clarifications for Diffserv*	IETF	2002

 © The Broadband Forum. All rights reserved

[29]	RFC 3315	*Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*	IETF	2003
[30]	RFC 3931	*Layer Two Tunneling Protocol - Version 3 (L2TPv3)*	IETF	2005
[31]	RFC 4271	*A Border Gateway Protocol 4 (BGP-4)*	IETF	2006
[32]	RFC 4719	*Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*	IETF	2006
[33]	RFC 4861	*Neighbor Discovery for IP version 6*	IETF	2007
[34]	RFC 5280	*Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*	IETF	2008
[35]	RFC 5340	*OSPF for IPv6*	IETF	2008
[36]	RFC 5357	*A Two-Way Active Measurement Protocol (TWAMP)*	IETF	2008
[37]	RFC 5880	*Bidirectional Forwarding Detection (BFD)*	IETF	2010
[38]	RFC 7348	*Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*	IETF	2014
[39]	NFV 001 V1.1.1	*Network Functions Virtualisation (NFV); Use Cases*	ETSI	2013
[40]	NFV-INF 001 V1.1.1	*Network Functions Virtualisation (NFV); Infrastructure Overview*	ETSI	2015

2.3 Definitions

The following terminology is used throughout this Technical Report.

BG	Business Gateway.
pBG	Physical Business Gateway. The Customer Premises Equipment (CPE) located at the business customer premises that contains all hardware-dependent BG functions that must be performed at the customer premises. It may have a built-in Network Function Virtualization Infrastructure (NFVI).
pBG-LAN interface	Interface(s) on the pBG for connecting local devices (e.g. office, computers).

 © The Broadband Forum. All rights reserved

vBG	Virtual Business Gateway. A virtual entity located at the network and/or at the customer site, serving one or more pBG entities, supporting some network and service functions such as IP routing.
LSL	Logical Subscriber Link. A logical point to point L2 connection between the pBG and the vBG.
pBG-LSL interface	The logical interface on the pBG facing the vBG.
vBG-LSL interface	The logical interface on the vBG facing the pBG.
vBG-WAN interface	Logical interface(s) on the vBG to one or more IP networks.
VBG System	A system that includes the pBG component at the customer site and the vBG component at the network and/or at the customer site. It also includes their connection over the LSL as well as their interfaces and management system.
NFVI-Node	As per ETSI GS NFV-INF 001 V1.1.1 [40], physical device deployed and managed as a single entity providing the NFVI functions required to support the execution environment for Virtualized Network Functions (VNFs).

2.4 Abbreviations

This Technical Report uses the following abbreviations:

3DES	Triple DES
ACL	Access Control List
ACS	Auto-Configuration Server
AES	Advanced Encryption Standard
ALG	Application Layer Gateway
AP	Access Point
ARP	Address Resolution Protocol
BGP	Border Gateway Protocol
BNG	Broadband Network Gateway
BOOTP	Bootstrap Protocol
BPCF	Broadband Policy Control Framework
CAPEX	Capital Expenditure
CHAP	Challenge-Handshake Authentication Protocol

 © The Broadband Forum. All rights reserved

CoA	Change of Authorization
CPE	Customer Premises Equipment
CPU	Central Processing Unit
DC	Data Center
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
DoS	Denial of Service
DPI	Deep Packet Inspection
DSCP	Differentiated Services Code Point
DSL	Digital Subscriber Line
EAP	Extensible Authentication Protocol
EMS	Element Management System
ETSI	European Telecommunications Standards Institute
EVC	Ethernet Virtual Circuit
FQDN	Fully Qualified Domain Name
GPON	Gigabit-capable PON
GRE	Generic Routing Encapsulation
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
IS-IS	Intermediate System to Intermediate System
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardization Sector
LAN	Local Area Network
MAC	Media Access Control
MD5	Message Digest 5
MPLS	Multi-Protocol Label Switching
MS-BNG	Multi-Service BNG
MSBN	Multi-Service Broadband Network
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NAPT	Network Address and Port Translation
ND	Neighbor Discovery
NFVI	Network Function Virtualization Infrastructure
NMS	Network Management System
NTP	Network Time Protocol

 © The Broadband Forum. All rights reserved

NUD	Neighbor Unreachability Detection
OAM	Operations, Administration and Maintenance
OPEX	Operational Expenditure
OS	Operating System
OSPF	Open Shortest Path First
OSS	Operational Support System
PAP	Password Authentication Protocol
PAT	Port Address Translation
PCP	Port Control Protocol
PE	Provider Edge router
PEAP	Protected EAP
PEP	Policy Enforcement Point
PGW	Packet Data Network Gateway
PON	Passive Optical Network
PNF	Physical Network Function
RADIUS	Remote Authentication Dial-In User Service
RED	Random Early Discard
SBC	Session Border Controller
SDN	Software Defined Networking
SHA	Secure Hash Algorithm
SLA	Service Level Agreement
SLAAC	Stateless Address Auto-configuration
SME	Small and Medium Enterprise
SNMP	Simple Network Management Protocol
SOHO	Small Office and Home Office
SP	Service Provider
SSID	Service Set Identifier
TCP	Transmission Control Protocol
TMF	Traffic Management Function
TR	Technical Report
TWAMP	Two-Way Active Measurement Protocol
TWL	TWAMP Light
UDP	User Datagram Protocol
VAS	Value-Added Services
vCPE	Virtual CPE
vE-CPE	Virtual Enterprise CPE
VID	VLAN Identifier
VLAN	Virtual LAN

 © The Broadband Forum. All rights reserved

VNF	Virtualized Network Function
VNFC	VNF Component
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VXLAN	Virtual eXtended LAN
WA	Work Area
WAN	Wide Area Network
WLAN	Wireless LAN
WRED	Weighted RED
WRR	Weighted Round Robin

© The Broadband Forum. All rights reserved

3 Technical Report Impact

3.1 Energy Efficiency

The VBG System architecture relocates some BG functionalities in the network. This will result in some increased consumption of energy in the service provider network. This can be mitigated by intelligent resource consumption mechanisms, whereby resources such as a firewall or WAN accelerator, etc., for a customer are turned off when not in use.

3.2 Security

TR-328 does have an impact on security. Various aspects of a single dedicated system with a routed BG 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 the VBG System will utilize shared resources, the vulnerability to Denial of Service (DoS) attacks in particular may increase. Moreover, the pBG may also provide NFVI capabilities, providing further opportunities for attacks. The provider of the VBG System services will need to address these issues.

3.3 Privacy

The move of some of the BG functionalities to the network may result in further exposure of the branch devices to the operator's network – they are no longer hidden behind an on-premise Network Address Translation (NAT) function – which may lead to some privacy concerns. This section compares legacy BGs to the VBG System architecture with regards to privacy.

3.3.1 Privacy with a routed BG

A managed BG is controlled by the Service Provider. In this model, the operator operates the BG and has access to detailed information about the status of the branch network and its devices. The Service Provider can configure the BG, collect statistics, etc. and, depending on the contract, could even activate some level of Deep Packet Inspection (DPI) in the BG or even mirror some Local Area Network (LAN) traffic.

A non-managed BG is not controlled by the operator. However, this does not prevent the operator from monitoring the outgoing traffic to the Internet and/or Business Virtual Private Networks (VPNs) at the network edge.

3.3.2 Privacy with a pBG in the VBG System architecture

3.3.2.1 Device visibility

LAN devices connected to a pBG have the same exposure to the operator as those connected to a managed legacy BG. However, with the VBG System architecture, some of the BG functions, including the firewall, are now located in the network. Therefore, the vBG function should be hosted in secure premises such as a PoP or Datacenter.

 © The Broadband Forum. All rights reserved

3.3.2.2 Device accessibility

Just as in the legacy architecture, access to LAN devices from the Internet and from other residential and business subscribers is prevented by the vBG firewall functionality.

Those VBG System subscribers concerned with privacy can insert their own router behind the pBG to hide their devices, similar to the BG case. However, this would prevent the subscriber accessing most of the VBG System specific services.

3.3.2.3 Local traffic

Traffic between two LAN devices located in the branch remains local, irrespective of whether the pBG is functioning in Bridged, Routed or Routing & Bridging modes. Only unicast traffic to be sent to the Internet, business VPN services, as well as broadcast and multicast messages for Bridged pBG's such as Dynamic Host Configuration Protocol (DHCP), Address Resolution Protocol (ARP), etc., are forwarded to the vBG. Local unicast traffic is not sent to the network, e.g. file transfer between two computers, network printing, etc.

3.3.2.4 Conclusion

The VBG System architecture is similar to the legacy managed BG architecture with respect to privacy. However, as routing and firewall functions are moved to the network, it is important that the vBG function be located in secure premises. Wary subscribers can add their own router behind their pBG to hide all or a subset of their devices.

 © The Broadband Forum. All rights reserved

4 Introduction

4.1 Business drivers

Service Providers communication services to enterprise customers have been evolving from commoditized site-to-site connectivity to value-added services, in order to move up the value chain. Typically, such new functionality is provided by adding a physical appliance at the customer site, providing services such as Firewall, Intrusion Detection System (IDS), WAN optimization, Session Border Controller (SBC), etc. Such solutions require multiple truck rolls and lead to a significant CAPEX/OPEX increase on upgrade/refresh, troubleshooting and/or maintenance operations.

The Network Functions Virtualization (NFV) direction, empowered with Software Defined Networking (SDN) technologies, enables introducing such new capabilities in a much more flexible and effective way - with VNFs that can be installed and service-chained on demand. As defined in the ETSI GS NFV 001; Use Cases [39], such virtualized functionality for virtualized enterprise CPE (vE-CPE or vCPE) can be placed in a central/distributed Data Center (DC), over NFV infrastructure, e.g. Cloud CO. Alternatively, Network Functions could be deployed at the customer site (the decentralized case), or both cases can be combined (the distributed case).

The emerging vCPE solution is driven by two major business needs: increasing revenues and optimizing costs.

Increasing revenues can be achieved by means of:

- **Reduction of Time-to-Market:** Achieved by introduction of new capabilities and services "at the speed of software" without changing customer-located equipment (no installation or upgrading of physical appliances). Multiple functionalities can be located in a single device using virtualization and service chaining.
- **New services and capabilities:** Allowing new sales opportunities (try-and-buy, pay-as-you-go, etc.), higher service variety and finer granularity, as well as upsell opportunities facilitated by the dynamic nature of the service.

Cost optimization will focus on:

- **CapEx reduction:** Obtained by using less expensive, simplified CPEs (in some cases), lower cost virtual appliances as opposed to physical, re-use of virtual appliance licenses and new service introduction.
- **OpEx reduction:** Will be made possible by network and service automation, decreased CPE software management complexity, in particular with network-based functionality and reduced truck rolls for maintenance and service upgrades.

 © The Broadband Forum. All rights reserved

4.2 Current Business Gateway deployment practices

The business environment can be composed of one site or multiple geographically separated sites, which need to be securely interconnected.

The Business Gateway is a physical equipment which is carrier owned and customer located. The BG is located at a business customer's site for the purpose of delivering carrier managed IP connectivity services such as IP VPN and Internet Access. It is traditionally implemented using a Customer Equipment (CE) router which performs such functions as WAN routing and restoration, protocol conversion/adaptation, tunneling, authentication, Network Address and Port Translation (NAPT), Access Control Lists (ACLs), QoS, and encryption. It may also provide LAN routing functions such as DHCP Server and inter-subnet routing. The BG allows the operator to deploy, control, manage and monitor the service being delivered to the customer premises.

The diagram, below illustrates the general deployment practice for the BG in the carrier's network:

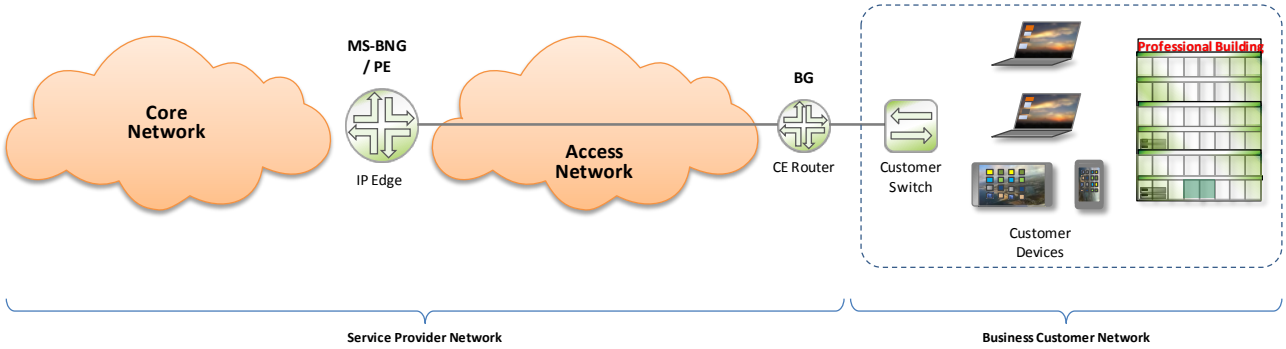

Figure 1 – Existing simple BG deployment

The service provider may choose to implement any of several protection mechanisms for the access tail into the business customer site, including logical line redundancy, e.g. using dual Ethernet Virtual Circuits (EVCs), physical line redundancy and even dual BG at the customer site.

In some cases, service providers will extend the functionality of the BG to offer additional managed services to the business. This is done through the deployment of specialized network appliances at the customer premise or in the network. For example: stateful firewalls, WAN optimizers, and Wireless LAN (WLAN) controllers. However, in many cases this is not cost effective due to the increased deployment and operational costs.

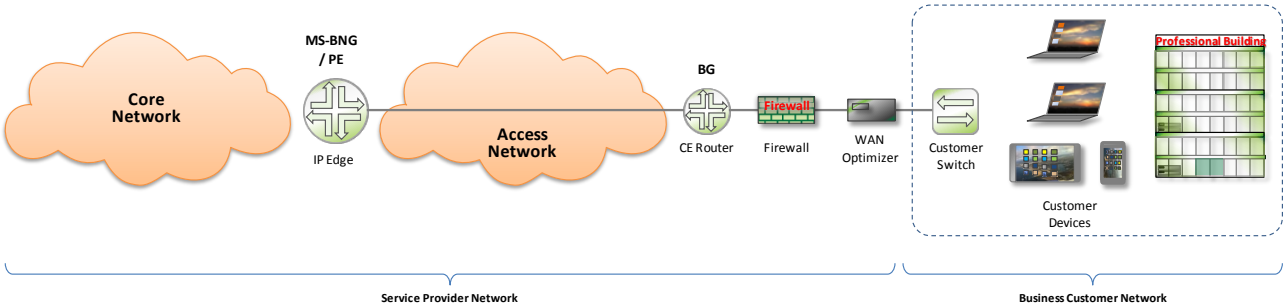

Figure 2 – Value-added services in current BG deployments

© The Broadband Forum. All rights reserved

4.3 Virtual Business Gateway (VBG) System overview

The Virtual Business Gateway System architecture is targeted at increasing the flexibility and service agility of the current Business Gateway, which hosts functions like DHCP, Firewall, NAT, IP forwarding, dynamic routing, etc. To do so, some of the traditional BG functions are moved to the network to enable more flexibility and eliminate the constraints introduced by the hardware in the customer site.

4.3.1 VBG System components

As result of the decomposition of the Business Gateway, its functionality is split into multiple components, as depicted in the following diagram:

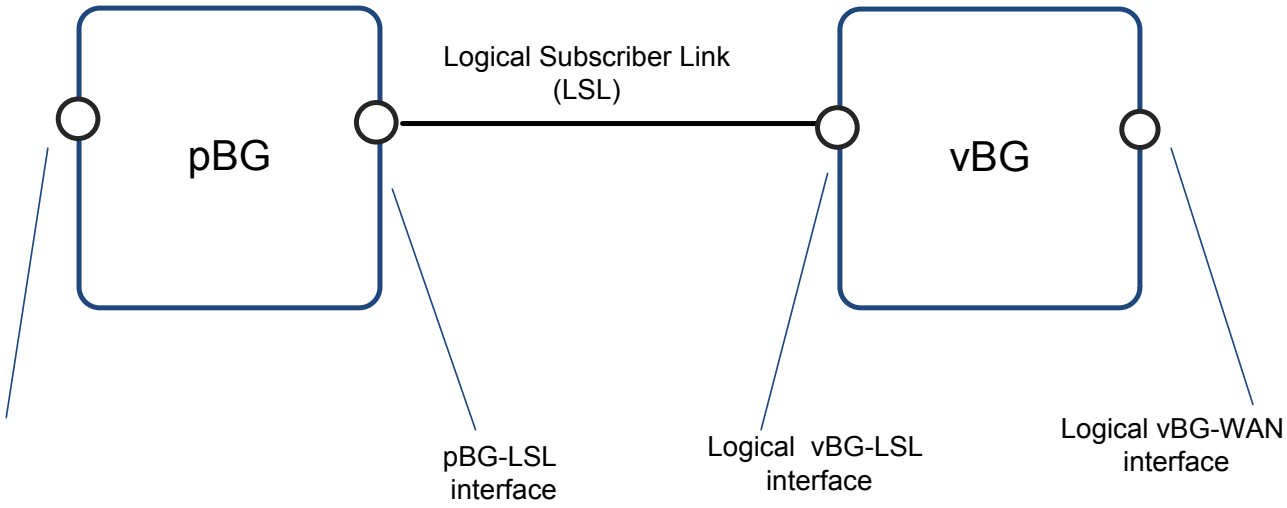

Figure 3 – A high level view of the VBG System

The pBG is the CPE located at the business customer premises and contains all the hardware-dependent BG functions that must be executed at the customer premises. The vBG is a self-contained or distributed virtual entity that performs the rest of the BG logical functions. For example, one may have a (hardware) traffic conditioning function at the customer premises performed by the pBG, while a (software) firewall function resides somewhere in the network, performed by the vBG.

The Logical Subscriber Link (LSL) is a logical point-to-point layer 2 connection between the pBG and the vBG, interconnecting the set of networking and service related Network Functions that compose the VBG System. The LSL has the following interfaces:

- pBG-LSL interface: Logical layer 2 interface on the pBG facing the vBG.
- vBG-LSL interface: Logical layer 2 interface on the vBG facing the pBG.

The pBG provides connectivity to one or more LAN broadcast domains by means of the pBG-LAN interface(s), and the vBG enables connectivity to one or more IP networks via the vBG-WAN logical interface.

© The Broadband Forum. All rights reserved

Note that the vBG logical functions may be distributed between the customer premises and the network. The functional split of BG functions between the pBG and the vBG will be discussed in the following sections.

4.3.2 vBG function location

The VBG System virtualizes some of the functions of BG, which is owned and provided by the service provider. These virtualized functions are contained within the vBG while the remaining physical functions are contained within the pBG. In this way, the service provider has additional flexibility in deploying new services within the vBG space without needing to install additional equipment at the customer premise.

The pBG provides functions which can only be provided at the customer premise. In addition, the pBG may provide NFV infrastructure for cases where some or all of the vBG functions are hosted at customer premises.

The service provider has flexibility to locate the vBG functions where it will provide an optimum combination of cost, performance and contractual obligations. The following sections describe some vBG placement examples.

4.3.2.1 vBG in the network

All vBG functions are placed in the network leaving only one physical device at the customer site, providing physical connectivity (e.g. LAN, WLAN), local Ethernet switching and transport to the vBG functions in the network, simplifying the CPE deployment.

The vBG functions may be hosted in the Multi Service Broadband Network Gateway (MS-BNG), in a highly scalable virtualization environment or a combination of both, e.g. L3, Routing, DHCP, NAT, etc. in the MS-BNG and Value-Added Services (VAS) functions such as IDS/IPS, WAN acceleration, etc. in VNFs. In the latter case, LSL termination happens at the MS-BNG.

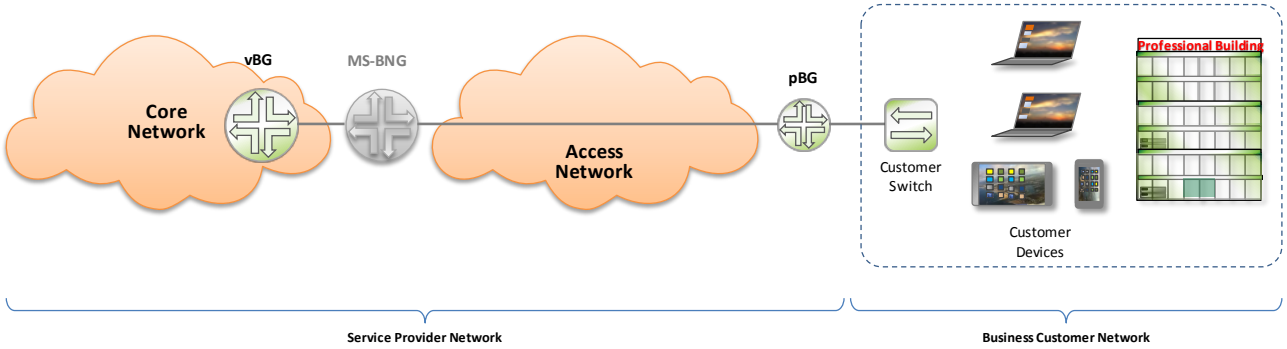

Figure 4 – VBG System deployment with vBG located in the network

© The Broadband Forum. All rights reserved

4.3.2.2 vBG at the customer premises

All vBG functions are placed at the customer premise requiring the pBG at the customer site to provide the NFVI. This allows the service provider to deploy the VBG System architecture without any need to change the existing network infrastructure or configuration, while adding some complexity in managing the compute nodes.

Added complexity will come from managing orchestration and deployment over a very large number of compute hosts, which are distributed everywhere in the service provider network. In addition, tightened security will be required in the NFV environment, due to the location of the compute hosts.

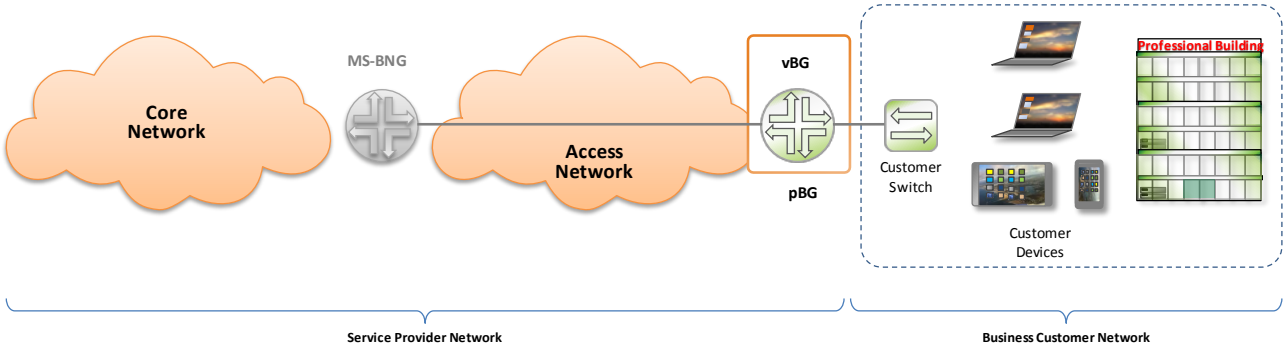

Figure 5 – VBG System deployment with vBG located at the customer premises

4.3.2.3 vBG distributed between network and customer premises

The vBG functions are distributed between the customer premise and the network. This may allow a smaller capacity NFVI to be provided in the pBG where only functions which must be provided at the customer premises can be hosted. This could be for reasons of security, performance / network topology or end user preference.

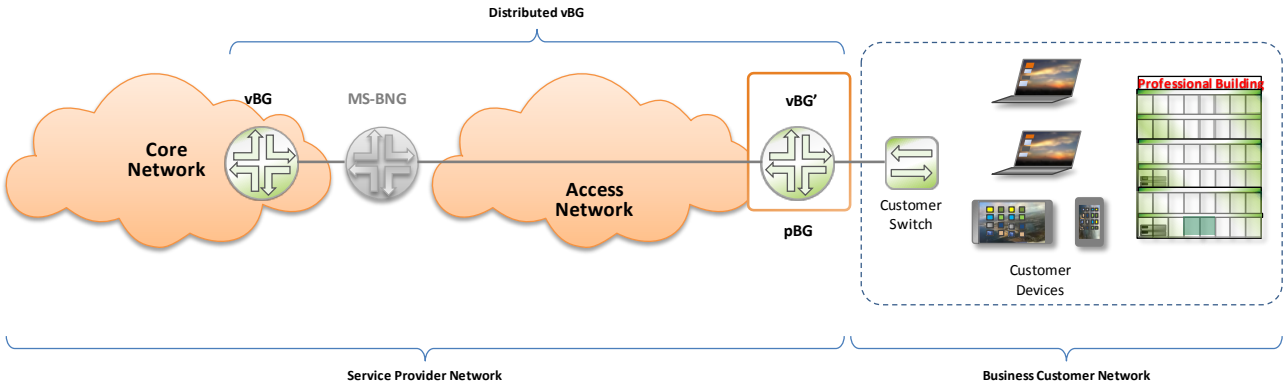

Figure 6 – VBG System deployment with distributed vBG

© The Broadband Forum. All rights reserved

4.3.3 Virtualization of compute resources

As depicted in Section 4/TR-359 [9] and section 4.3.4 in this document, a VBG System may consist of PNFs (Physical Network Function) and/or VNFs (Virtual Network Functions) located in the NFVI, at the customer site and/or in the network. The VBG System, specially the vBG function, could be "decomposed" in various VNFs or VNF Components (VNFCs) running in one or multiple NFVI PoPs. Hence the architecture should facilitate the decomposition and relationship between VNFs defined in ETSI GS NFV-INF 001 V1.1.1.

4.3.4 vBG hosting infrastructure

The vBG Hosting Infrastructure:

- Hosts the vBG functions, which can be either VNFs or Physical Network Functions (PNFs)
- Terminates the LSLs connecting the pBG's to their respective vBG's and connects the vBG's to the IP network

This platform may consist of one single network element, such as a MS-BNG or pBG, or may be composed of a set of network elements, such as a pBG, MS-BNG and NFV infrastructure (COTS, ToR, etc.).

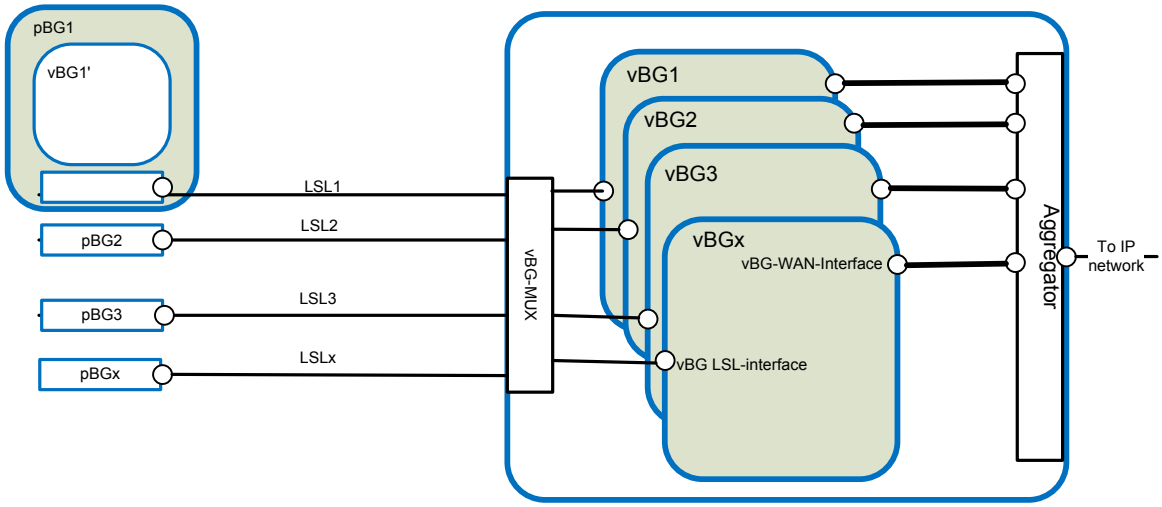

Figure 7 – Flexible vBG hosting infrastructure

Figure 7 depicts a case where the vBG is located only at the Service Provider network, illustrated by vBG's 2,3 and N, and a case where the vBG is distributed between the Service Provider network and the customer site (pBG), represented by vBG 1. In this case, the LSL is still terminated in the pBG even though vBG1' is behind it. vBG1' hosts the Network Functions that must be located at the customer premises.

The vBG_MUX (vBG Multiplexer) is the entry point of the vBG hosting infrastructure. It is responsible for mapping an enterprise branch to a vBG. The LSL that connects a pBG to its vBG is composed of two segments: the first segment connects the pBG to the vBG_MUX. The second segment extends the LSL from the vBG_MUX to the appropriate vBG in the vBG hosting infrastructure.

© The Broadband Forum. All rights reserved

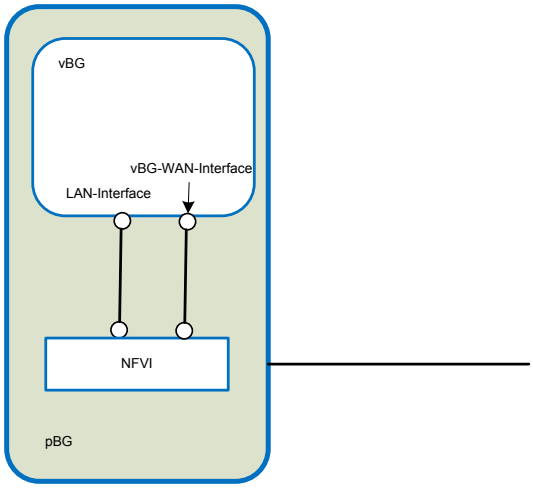

Figure 8 – vBG hosting infrastructure at the customer premises

In Figure 8, the full vBG and pBG reside at the Customer site and, in this case, there is no need for LSL. The vBG WAN interface will be connected to pBG forwarding entity.

4.4 High level architectural components and functional distribution overview

Figure 3 describes the pBG and vBG as components of the VBG System, to which Network Functions are distributed. Following are the VBG System requirements:

- LAN Interfaces (e.g., Ethernet, Wi-Fi, Voice Analog Telephone Adapter (ATA))
- WAN Interface(s)
- LSL Interface
- Learning Bridge
- IPv4/IPv6 Forwarding
- Routing Control Plane
- LAN QoS
- Upstream QoS
- Downstream QoS
- WAN QoS
- DHCP Client
- IP Address Management / DHCP Server
- Network Address and Port Translation (NAPT)
- DNS
- OAM
- Performance Monitoring
- Security
- Management Client
- Control Client
- Authentication Client
- NTP Client

 © The Broadband Forum. All rights reserved

These capabilities, which are distributed to the pBG and/or vBG, are depicted in Figure 9 based on the high level architecture in Figure 3.

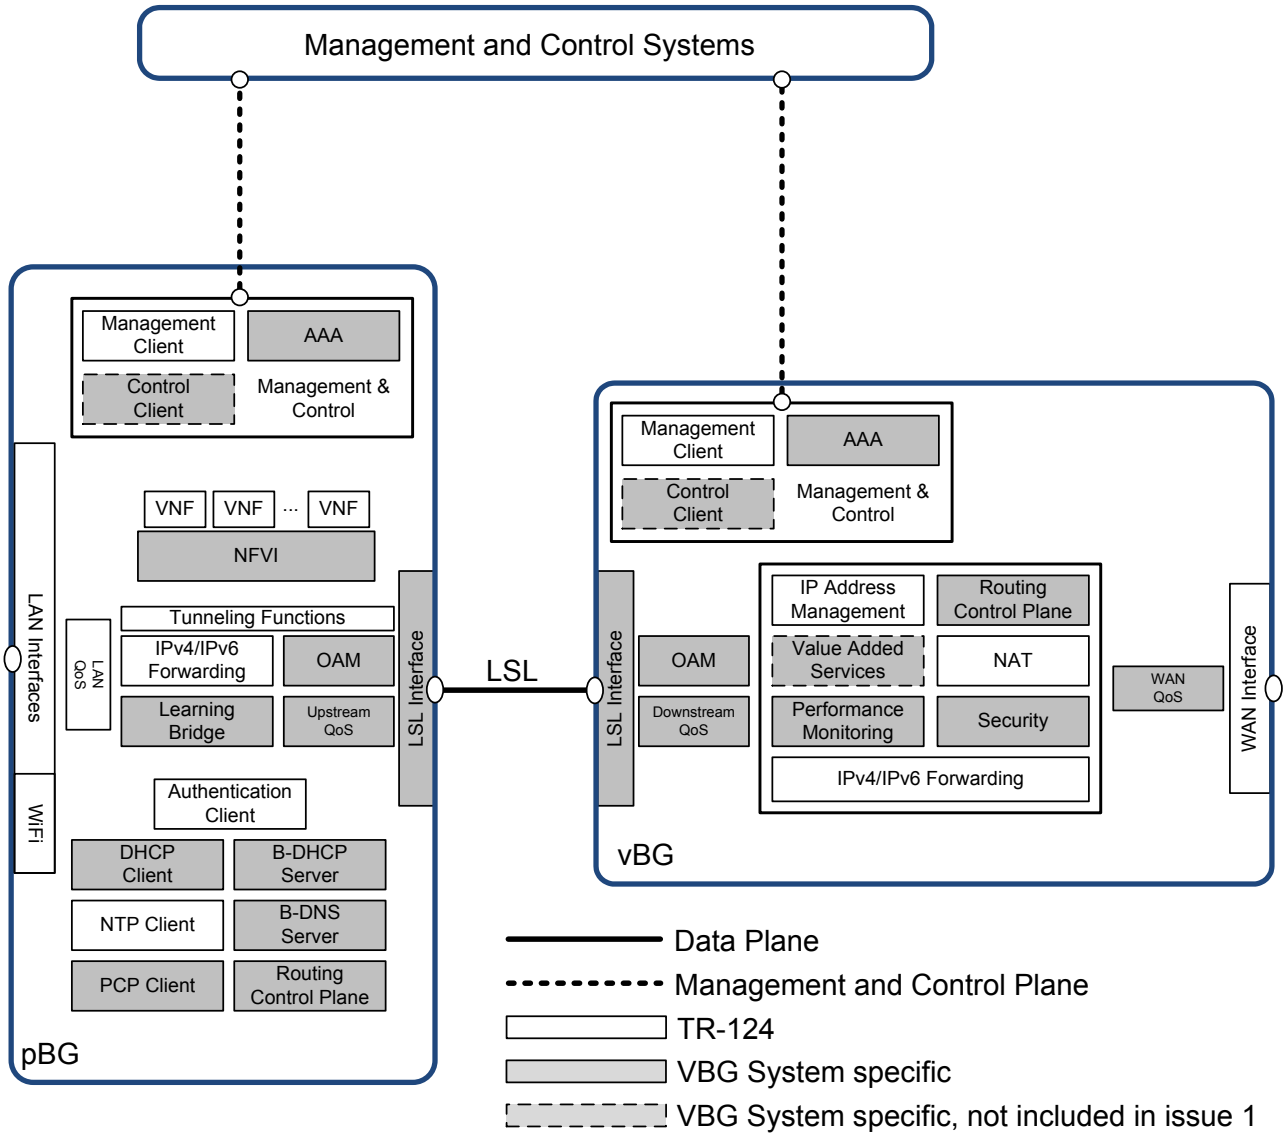

Figure 9 – Functional Distribution of VBG System

The following capabilities go beyond TR-124 [3] and are specific to the VBG System, either to support the connectivity between the pBG and vBG or to enable VBG System specific use cases:

- LSL Interface – section 6.1
- LSL Monitoring and Protection – section 6.3
- Routing Control Plane – section 6.7
- Value Added Services – out of scope of the current issue of TR-328

Note: The Functional Distribution does not intend to prescribe a preferred implementation or system design but to define requirements associated with the identified VBG System.

© The Broadband Forum. All rights reserved

4.4.1 pBG functional architecture

The pBG provides Network Functions to forward data between devices in the subscriber's premises and Network Functions of the vBG, for vBG functions that are located in the customer site. In that case, the pBG includes an NFVI-Node. In addition, the pBG enables management of the pBG itself and its NFVI-Node functionality, forwards management traffic of its vBG hosted functions and provides capabilities for the troubleshooting of the LSL. Finally, the pBG contributes to increasing the reliability of the VBG System by supporting redundancy, including multi-homing.

As indicated above two types of pBG can be found:

1. pBG without NFVI
2. pBG with built-in NFVI (it enables running virtual functions at the customer site)

The following lists capabilities of the VBG System, applicable to the pBG:

- LAN Interfaces – section 4.5.1
- LSL Interface – section 6.1
- DHCP Client – section 6.1.4
- OAM – section 6.3
- Performance Monitoring – section 6.4
- LAN QoS – section 6.5
- Upstream QoS – section 6.5
- Learning Bridge – section 6.7.1
- IPv4/IPv6 Forwarding – section 6.7.2
- Routing Control Plane – section 6.7.4
- Authentication Client – section 6.9.2.1
- Management Client – section 6.11.1
- Control Client – out of scope of the current issue of TR-328
- NTP Client – section 7.4.5/TR-317 [8]
- In case line redundancy is not active:
    a. B-DHCP Server – section 7.4.2.3/TR-317
    b. B-DNS – section 7.4.2.4/TR-317

The use of these capabilities will depend on the pBG forwarding mode.

 © The Broadband Forum. All rights reserved

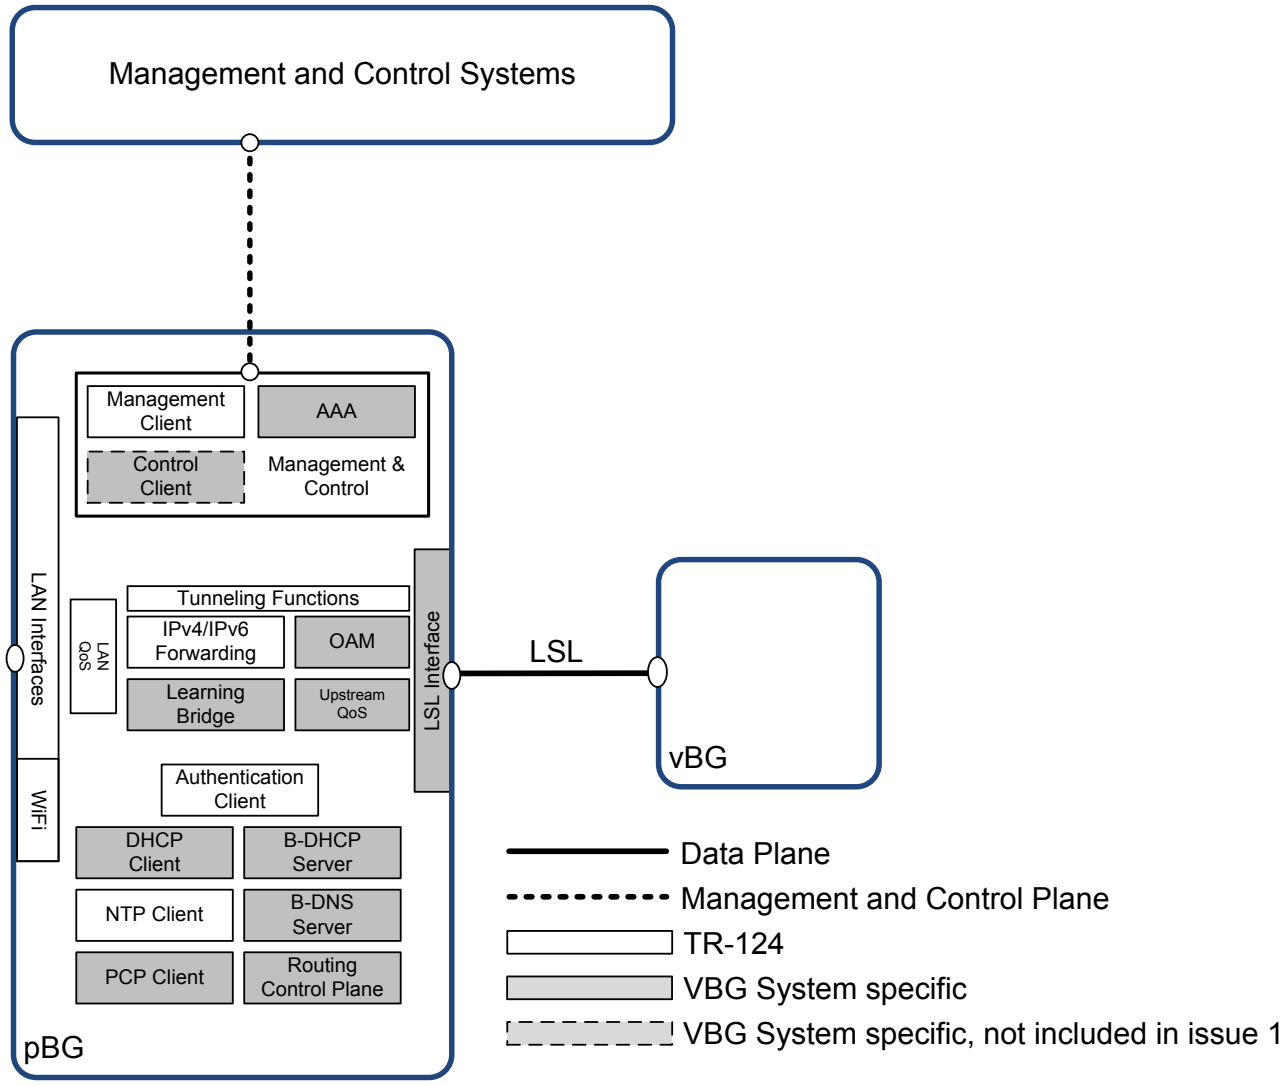

Figure 10 – pBG without built-in NFVI-Node functional architecture

For pBG with built-in NFVI-node, the following additional capabilities are required:

- pBG virtualization software: Operating System, hypervisor or container infrastructure, etc.
- pBG software and/or hardware: Switching/forwarding interface between the pBG-hosted functions and the LAN/WAN interfaces.

© The Broadband Forum. All rights reserved

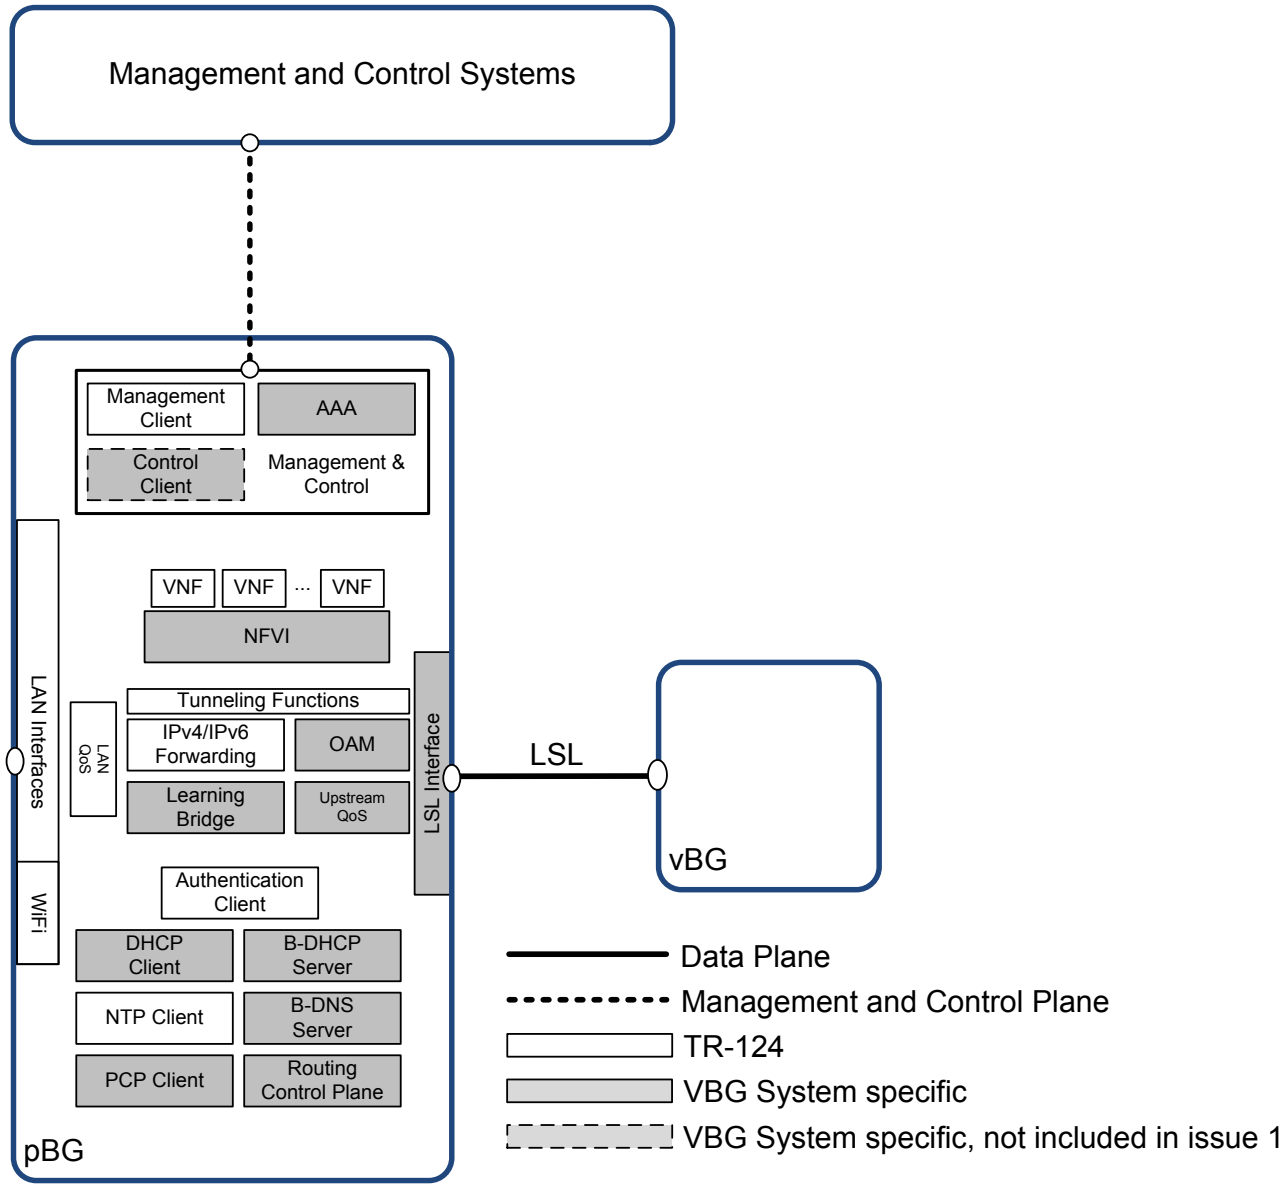

Figure 11 – pBG with built-in NFVI-Node functional architecture

4.4.2 vBG functional architecture

The vBG provides Network Functions that encompass the VBG System capabilities that may be distributed to the Service Provider network and/or NFVI-Node hosted at the pBG.

The following lists capabilities of the VBG System, applicable to the vBG:

- WAN Interface
- LSL Interface – section 6.1
- OAM – section 6.3
- Performance Monitoring – section 6.4
- Downstream QoS – section 6.5
- WAN QoS – section 6.5

 © The Broadband Forum. All rights reserved

- IP Address Management / DHCP Server – section 6.6.1
- Network Address and Port Translation (NAPT) – section 6.6.2
- IPv4/IPv6 Forwarding – section 6.7.3
- Routing Control Plane – section 6.7.5
- Security – section 6.8
- AAA / Authentication Client – section 6.9
- Management Client – section 6.11.2
- Control Client – out of scope of the current issue of TR-328
- Value Added Services – out of scope of the current issue of TR-328

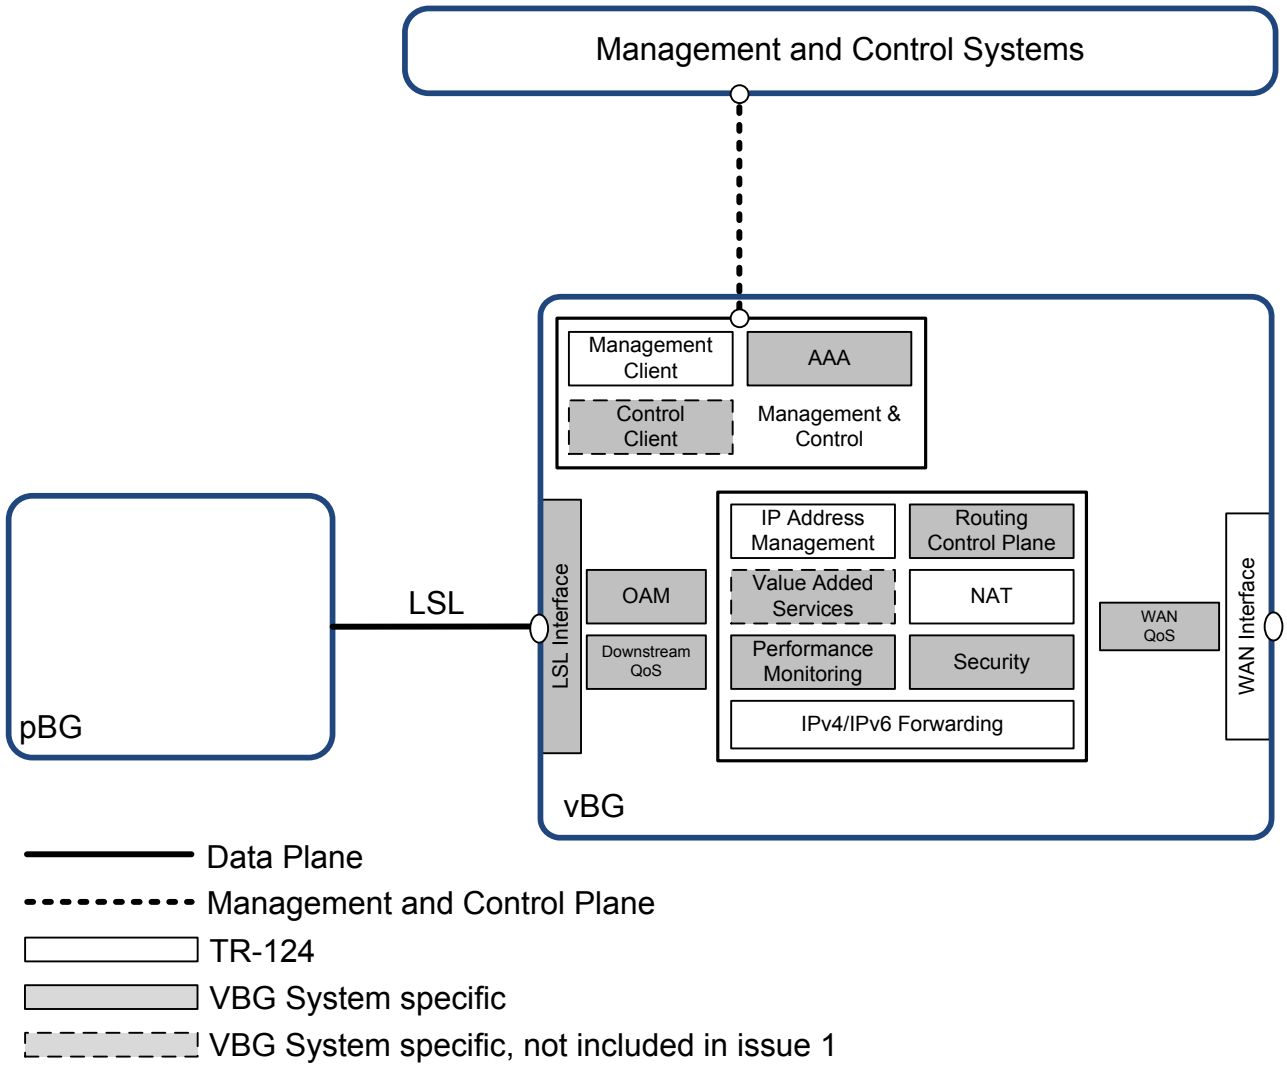

Figure 12 – vBG functional architecture

© The Broadband Forum. All rights reserved

4.5 Enterprise branch connectivity architecture

The business environment can be composed of one site or multiple geographically separated sites of different sizes, which need to be securely interconnected. This section presents architectural choices for the branch LAN and WAN connectivity.

4.5.1 Branch LAN

While the typical small branch will have a single Ethernet broadcast domain and a single IP subnet for all devices in the branch, a more complex business site may have multiple subnets for context separation (e.g. different departments) and/or security rule enforcement.

All devices for a given Ethernet broadcast domain may be connected directly to the pBG or the usual techniques may be used for Ethernet network extension, e.g. hubs/switches, Virtual LANs (VLANs), WLAN Access Points (APs), etc.

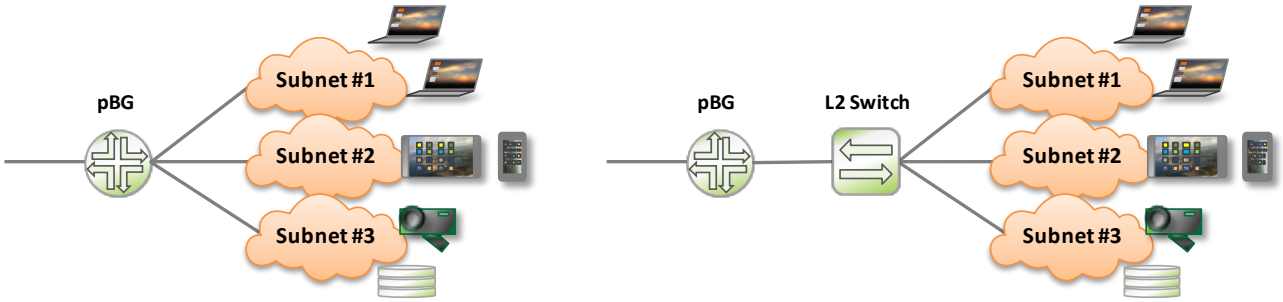

Figure 13 – Branch LAN – Devices connected directly or through L2 extension

In the cases where multiple IP subnets are required, inter-subnet forwarding may be achieved in different ways, depending on the pBG forwarding mode (see section 0).

A pBG routed model naturally solves the need for L3 forwarding between subnets. In order to enforce security constraints, the pBG L3 function will require filtering capabilities.

For a pBG bridged model, a separate routed CPE can be placed between the pBG and LAN devices that provides inter-subnet forwarding between the different branch subnets. This model has severe limitations as normally that CPE will not be controlled by the VBG System, and as such provisioning and operation of such device and security rule enforcement between subnets will be left to the enterprise.

© The Broadband Forum. All rights reserved

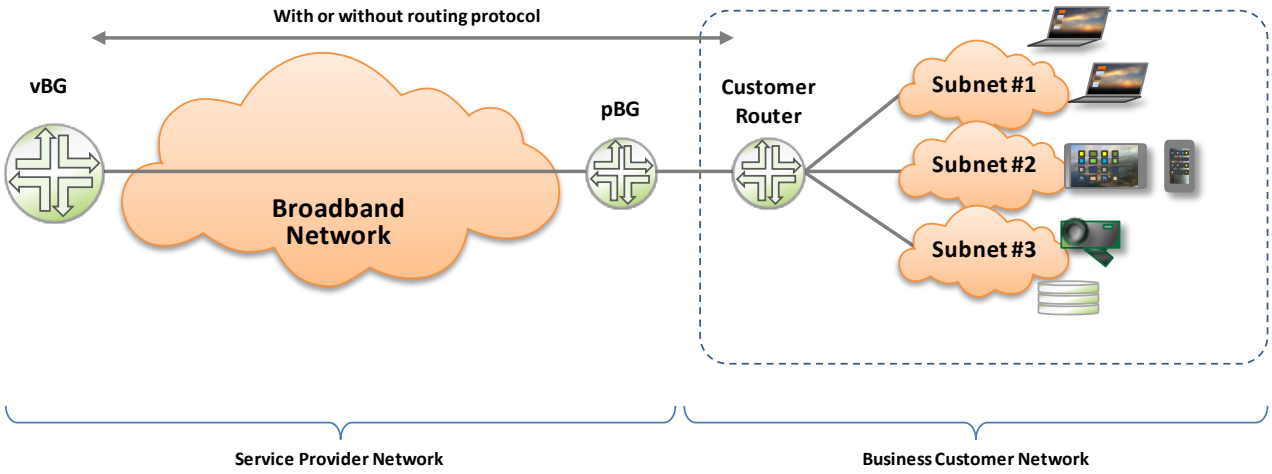

Figure 14 – Branch LAN – Inter-subnet forwarding with router behind pBG

A more suitable option for the case of a pBG bridged model is to forward all inter-subnet traffic to the vBG and have the vBG provide the L3 forwarding between subnets. This model requires extending each subnet from the LAN to the vBG, e.g. using separate VLANs for each subnet between pBG and vBG. In this model, the local traffic will take unnecessary bandwidth from the WAN and it can be solved by placing the Routing function of the vBG in the customer site.

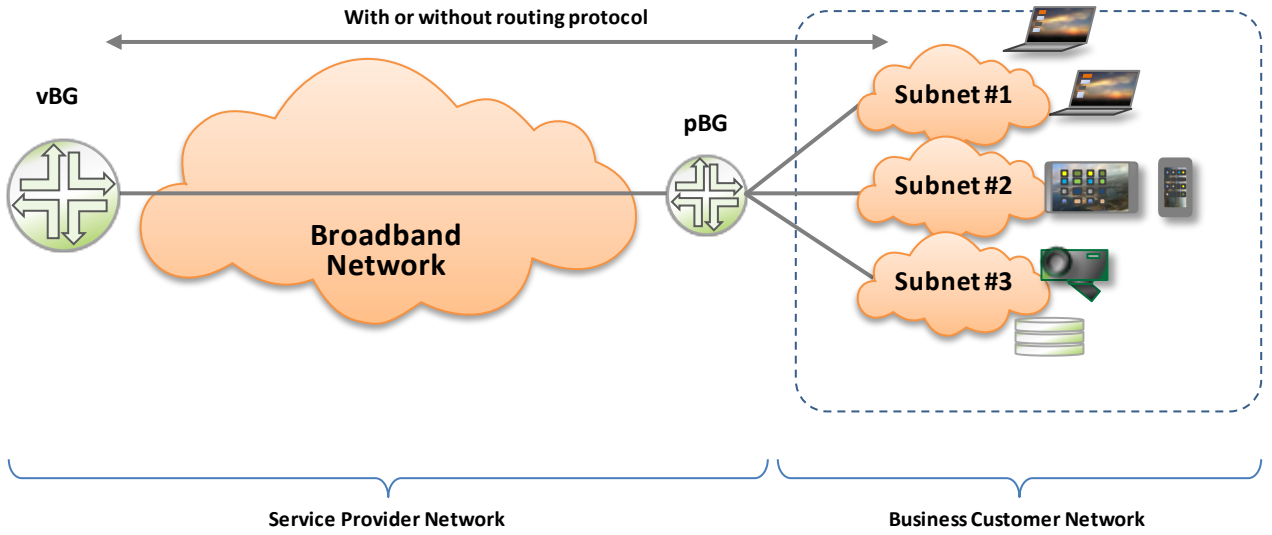

Figure 15 – Branch LAN – Inter-subnet forwarding by LAN extension to vBG

In the cases where inter-subnet forwarding is performed in the branch (routed pBG or routed CPE behind pBG), routes for each of the branch subnets will need to be set at the vBG, either based on RADIUS activation (e.g. Framed-Route) or using static or dynamic routing.

© The Broadband Forum. All rights reserved

4.5.2 Branch WAN

There are multiple aspects to consider in regard to WAN connectivity of a business branch site. The following is a list of some of those key aspects:

- pBG to vBG transport
- Uplink redundancy
- Network redundancy
- pBG redundancy: two or more redundant pBG in a business branch

pBG redundancy is out of scope of the current issue of TR-328.

4.5.2.1 pBG to vBG transport

In its most basic form, the connectivity between the pBG and the vBG can be Ethernet based, optionally transported over L2VPN services over the carrier network until reaching the vBG. This connectivity model is termed Flat LSL model in TR-328.

Note that in the case of a bridged pBG, and when transiting through a shared physical access network, using a dedicated VLAN (1:1) between the pBG and vBG will be required as a large number of LAN protocols have not been extended to identify the access line (e.g. DHCP Option 82) and as such a shared VLAN (N:1) between multiple pBG devices would cause issues.

In a number of cases plain Ethernet connectivity will not be sufficient. Such cases include transiting over any non-L2 network, such as a 3GPP mobile network or the case of servicing off-net sites, being aggregated over the Internet (a L3 network).

In such cases, a tunnel can be used between the pBG and the vBG providing connectivity over any IP network. This model is termed Overlay LSL model in TR-328.

4.5.2.2 Uplink redundancy

To increase the service availability to an enterprise branch site, more than one uplink can be installed in the pBG.

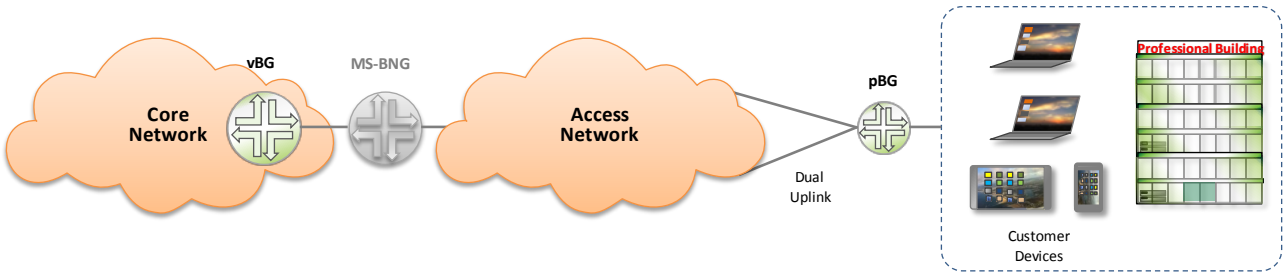

Figure 16 – Branch WAN – Uplink redundancy

Depending on the availability requirements, this may come in the form of a second physical or logical uplink. In the case of a second physical uplink, two physical ports and their respective access links (e.g. xDSL, PON) are used in the pBG. In contrast, in the case of a second logical

© The Broadband Forum. All rights reserved

uplink, the last mile access path is not redundant and only connectivity through the network is, e.g. two EVCs provided by the service provider from the pBG to the vBG.

4.5.2.3 Network redundancy

Reliability is one of the main concerns for business customers. In addition to uplink redundancy and if availability requirements mandate so, the pBG can be multi-homed to different access networks, e.g. xDSL/GPON and 3GPP mobile networks, to provide the branch with a more reliable connection to the Internet or to the other sites.

Moreover, the different access networks may be provided by different Service Providers to increase the reliability even further.

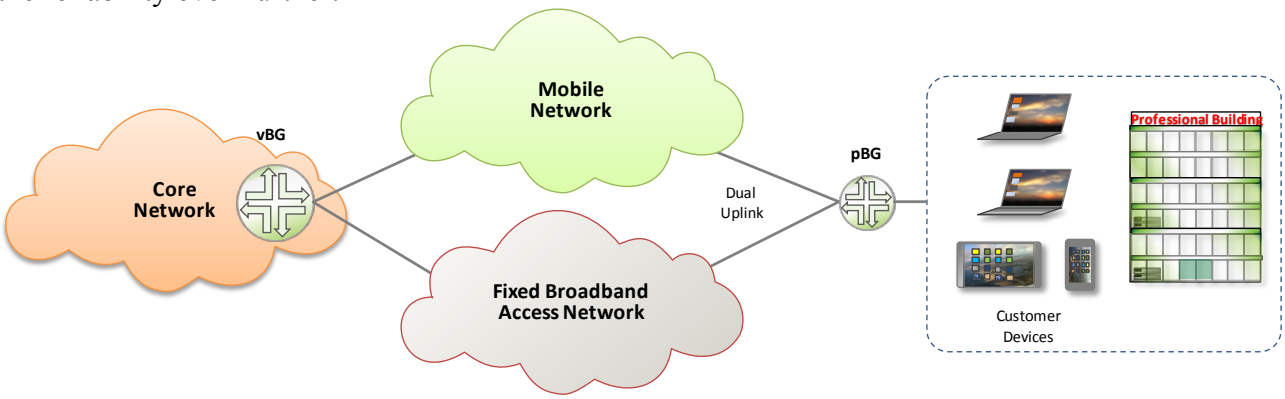

Figure 17 – Branch WAN – Network redundancy

The two uplink connections would typically be running on separate physical mediums - separate cables, fibers, radio interfaces or any combination of them.

For further details on redundancy scenarios, see section 6.2.

© The Broadband Forum. All rights reserved

4.5.3 pBG forwarding modes

The Virtual Business Gateway System is defined by the unbundling of features from customer located hardware. With this broad definition, there are many different deployment models, depending on which functions are extracted from the BG and where these functions are relocated to. In the same way that VPN services can have multiple models to answer different needs, the diversity of VBG System models adapts to a variety of diverse requirements.

TR-328 architectures support several pBG forwarding modes:

- Routed pBG: performs IP forwarding and routing
- Bridged pBG: performs L2 forwarding and MAC learning
- Routed & Bridged pBG

The following text categorizes the use cases and provides some guidelines as to what mode to be used in the pBG.

4.5.3.1 Case 1: Managed Router Service for SMB and SOHO

In general, Small and Medium-sized Business (SMB) and Small Office Home Office (SOHO) customers usually little to no technical expertise, so a simplified service will have preference over flexibility and feature richness.

The VBG System must support both IP VPN and Internet access services. It must provide the core features of a branch router, in particular addressing (DHCP server or relay), basic routing, access-lists, statistics, etc. The remaining on-site pBG must at least support L2 forwarding, QoS (L3 classification and uplink scheduling), L3 performance measurement and remote management. Optionally, it may have simple IP forwarding capabilities, in the case multiple local subnets are required in the branch.

4.5.3.2 Case 2: Managed Router Service for Medium and Large Enterprises

While similar to the managed router service for SMB, there are some significant differences, in particular in terms of technical expertise, complexity of the organization and reliability requirements. In addition to the SMB related features, the pBG may provide more complex connectivity and protection options and other features as follows:

- Local functions on site controlled by customer (e.g. DHCP server)
- Link redundancy / CE redundancy
- Complex routing (e.g. for multi-homing to various Service Providers)
- Large LAN infrastructure, including routers/switches/Aps/servers/etc.
- Multiple network domains / subnets

4.5.3.3 pBG forwarding mode selection criteria

The selection of what forwarding mode to use for the pBG depends on multiple factors. Among others:

 © The Broadband Forum. All rights reserved

- Connectivity to customer devices, e.g. number of LAN ports required, customer-owned routed device behind the pBG, etc.
- Service provider's network transport to the PoP/DC (L2, L3, Out of Franchise, etc.)
- Division of the network functions to virtual or physical
- Redundancy needs, e.g. redundant pBG uplinks, dual homing, dual pBG

The table below summarizes the recommended pBG forwarding mode based on the key criteria:

Number of subnets	Customer-owned device	LSL Type	pBG mode
Single subnet	None	Flat	Bridged
		Overlay	Bridged*
	Yes, switched	Flat	Bridged
		Overlay	Bridged*
	Yes, routed	Flat	Bridged
		Overlay	Bridged*
Multiple subnets	None	Flat	Routed
		Overlay	Routed
	Yes, switched	Flat	Routed
		Overlay	Routed
	Yes, routed	Flat	Bridged
		Overlay	Bridged*

Table 1 – Recommendations for pBG forwarding mode selection

Note: The cases requiring the use of Overlay LSL mode, with tunneling between the pBG and the vBG, will require L3 capabilities in the pBG. However, these can remain basic and limited to obtaining a WAN IP address (e.g. using DHCP or PPPoE) and being able to source a tunnel from that address. Traffic forwarded over the tunnel can then be L2 forwarded or routed depending on the selected pBG forwarding mode. In the redundancy case, more sophisticated mechanism might be used like routing protocols and in some cases policy based routing as well.

 © The Broadband Forum. All rights reserved

4.5.3.4 LSL encapsulation

The figures in this section depict the respective encapsulations for the various pBG forwarding modes and LSL types. It is important to note that the LSL can be native Ethernet with or without VLAN as described in section 6.1.1, or tunneled using the Overlay LSL tunneling options described in section 6.1.2.

4.5.3.4.1 Bridged pBG – Flat LSL

The figure below presents the case where the pBG is a bridge and the LSL is plain Ethernet with VLAN encapsulation. The MAC address in the Ethernet header (C-Eth) over the pBG to vBG link is the customer device's MAC address.

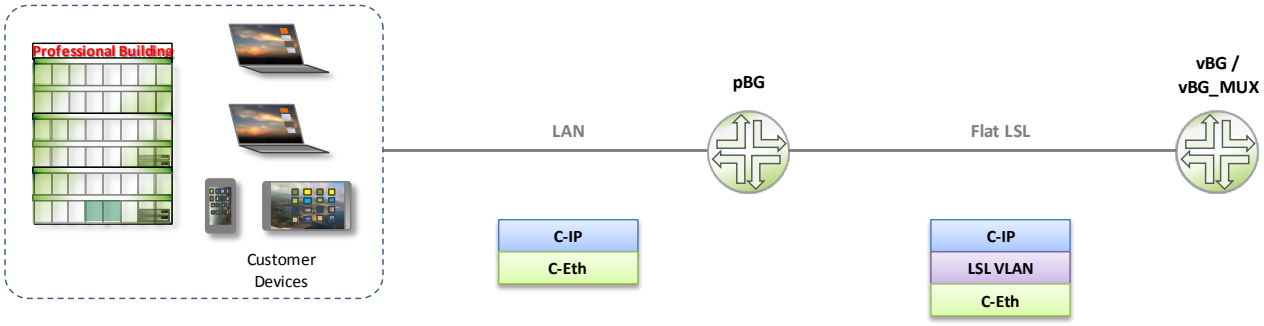

Figure 18 – Bridged pBG – Flat LSL

4.5.3.4.2 Bridged pBG – Overlay LSL

The figure below presents the case where the pBG is a bridge and the LSL is using one of the Overlay LSL tunneling options described in section 6.1.2. The MAC address in the outer Ethernet header (LSL Eth), in green, is the pBG uplink port MAC address. The MAC address in the inner Ethernet header (C-Eth) remains the customer device's MAC address.

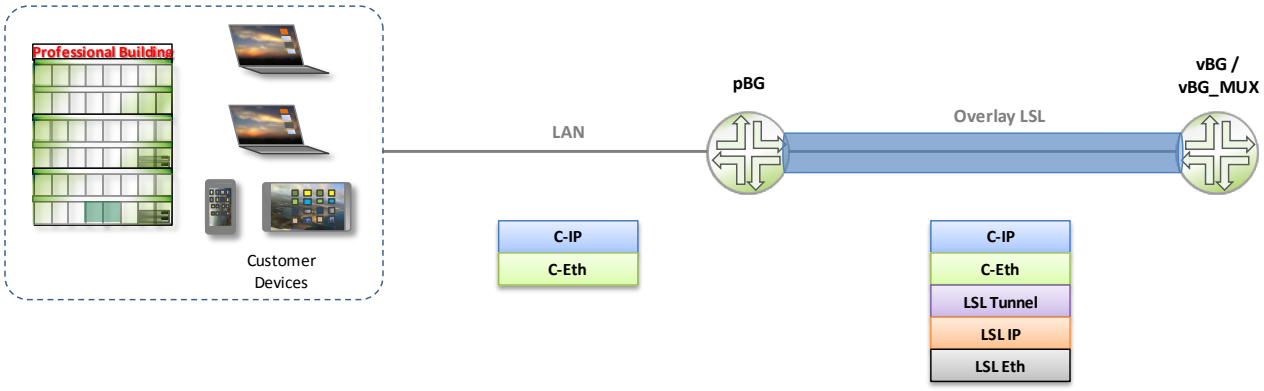

Figure 19 – Bridged pBG – Overlay LSL

© The Broadband Forum. All rights reserved

4.5.3.4.3 Routed pBG – Flat LSL

The figure below presents the case where the pBG is a router and the LSL is plain Ethernet with VLAN encapsulation. The MAC address in the Ethernet header (P-Eth) over the pBG to vBG link is the pBG uplink port MAC address.

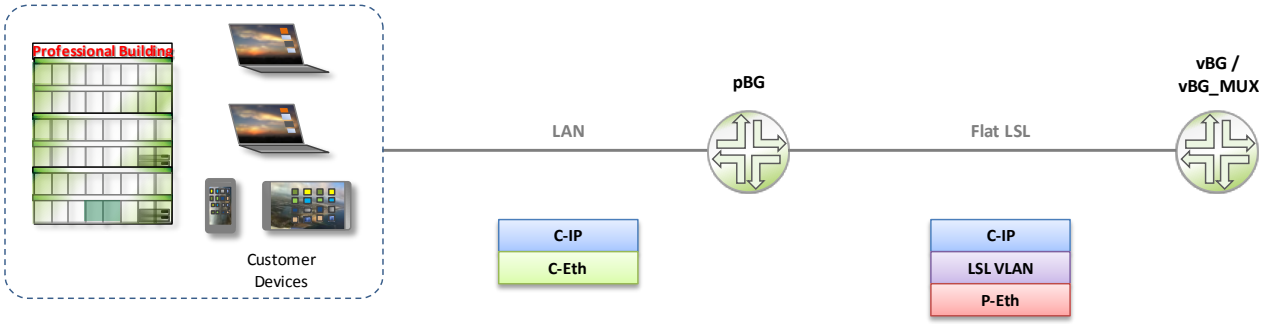

Figure 20 – Routed pBG – Flat LSL

4.5.3.4.4 Routed pBG – Overlay LSL

The figure below presents interface diagram where the pBG is a router and the LSL is using one of the Overlay LSL tunneling options described in section 6.1.2. The MAC address in the outer Ethernet header (LSL Eth), is the pBG uplink port MAC address. The MAC address in the inner Ethernet header (P-Eth) is also a pBG MAC address, which may or may not be the same as the pBG uplink port MAC address.

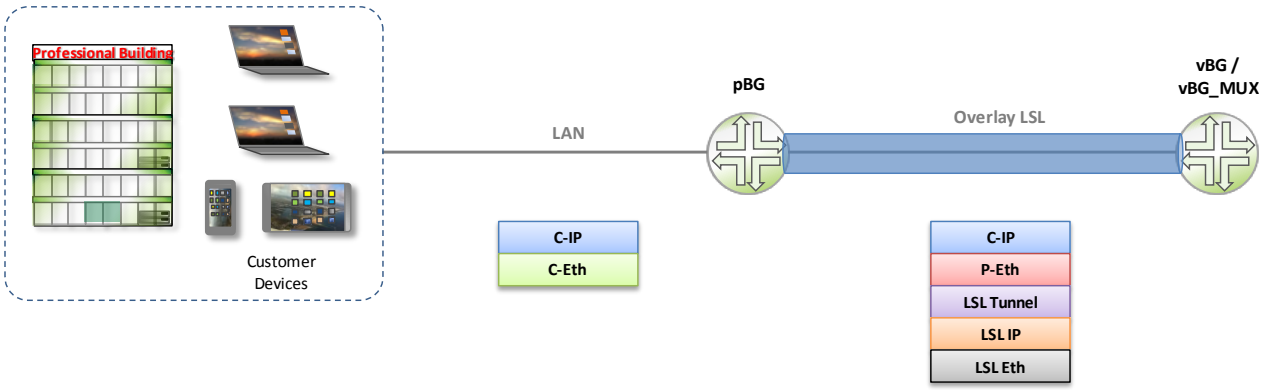

Figure 21 – Routed pBG – Overlay LSL

© The Broadband Forum. All rights reserved

5 VBG management and control

5.1 Management and control functional architecture

The network functions provided by the pBG and vBG are configured and maintained via one or more client functions associated with the pBG or vBG. These client functions are controlled by the associated management or controller system functions across a set of defined reference points.

Figure 4/TR-359 and Figure 5/TR-359 illustrate a combination of the BBF and ETSI-NFV architectures predicated on the assumption that services are hosted in NFVI that has augmented a BBF specified multi-service broadband network.

TR-359, the BBF Framework for Virtualization, defines the relevant reference points to the network functions. For the network functions of the pBG and vBG, these include the Minf, Ms, B and R reference points that are associated Management, Authentication, and Policy Client functions of the pBG and vBG.

Note: TR-359 does not include a reference point from the SDN Controller to the NF or NFVI. This specification adds a new reference point (RCI) that augments reference points in TR-359. Likewise, as described in section 4.3.3, TR-134 [4] does not describe a reference point for the Policy client in pBG. As such, this specification adds a new reference point (Rc).

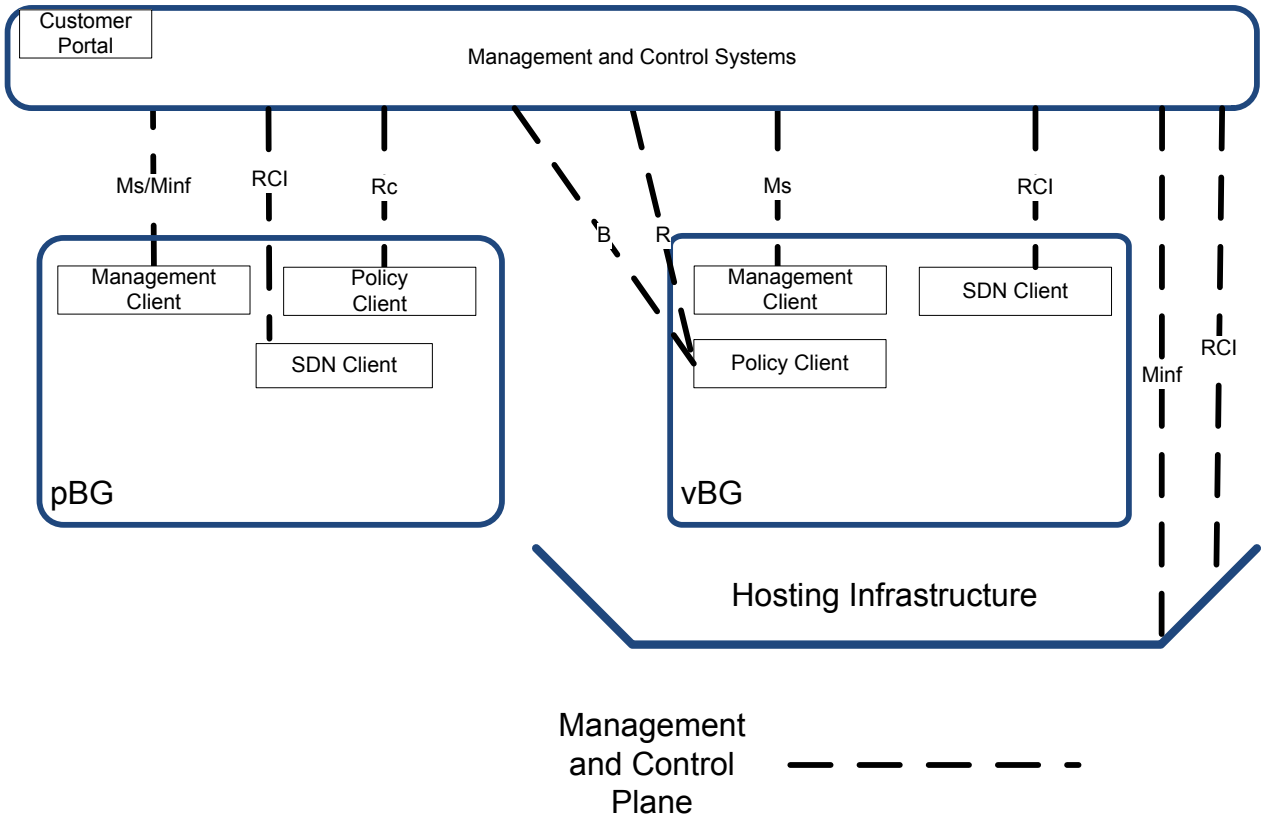

Figure 22 – Management and Control Architecture

© The Broadband Forum. All rights reserved

5.1.1 Application Functionality Management Reference Point (Ms)

The Ms reference point permits configuration and monitoring of the application functionality of network functions within the pBG and vBG, such as the DHCP Server, Firewall, etc. This reference point represents interfaces to either existing OSS/NMS/EMS management systems and provides FCAPS management to include:

- Fault Management: Alarm generation and troubleshooting (OAM) of network functions,
- Performance Management: Generation of performance measurements of network functions.
- Configuration Management: Configuration of the application layer functionality of network functions.
- Topology Management: Management of the topology of the nodes and network functions.
- Accounting: Usage measurements of the network functions.
- Security: Security associated with the application layer functionality of network functions.
- Software Management: Management of software and associated artifacts (configuration files) associated with network functions.

For management of the application layer functionality in the pBG, the following sections identify the requirements that need management across the reference point:

- Section 6.11.1 pBG Management Client

For management of the application layer functionality in the vBG, the following sections identify the requirements that need management across the reference point:

- Section 6.11.2 vBG Management Client

5.1.2 Infrastructure Management Reference Point (Minf)

The Minf reference point permits configuration and monitoring of the non-NFVI components of the pBG host and vBG hosting infrastructure (e.g., network interfaces, memory, CPU, OS, storage). This reference point represents interfaces to either existing OSS/NMS/EMS management systems and provides FCAPS management to include:

- Fault Management: Alarm generation and troubleshooting (OAM) of the host,
- Performance Management: Generation of performance measurements of the host
- Configuration Management: Configuration of the application layer functionality of the host.
- Security: Security of the host.
- Software Management: Management of OS software and associated artifacts (configuration files) associated with the host.

For management of the pBG infrastructure and networking functionality, the following sections identify the requirements across the reference point that needs management:

- Section 6.1 End-to-end network requirements
- Section 6.2 Multi-homing
- Section 6.3.1.1 Connectivity monitoring by the pBG
- Section 6.3.2 LSL Failure and Protection
- Section 6.4 Performance Monitoring requirements
- Section 6.5.1 QoS requirements on the pBG

 © The Broadband Forum. All rights reserved

- Section 6.7 Forwarding and routing protocols
- Section 6.11.1 pBG Management Client

For management of vBG Hosting Infrastructure, the following sections identify the requirements across the reference point that needs management:

- Section 6.1 End-to-end network requirements
- Section 6.3.1.2 Connectivity monitoring by the vBG and vBG_MUX
- Section 6.4 Performance Monitoring requirements
- Section 6.5.2 QoS requirements on the vBG
- Section 6.7 Forwarding and routing protocols
- Section 6.11.2 vBG Management Client

5.1.3 Policy Reference Points (B, R)

The B and R reference points permit the configuration of the vBG_MUX and vBG network functions acting in the role of Policy Enforcement Points (PEP) by different policy servers in the Multi-Service Broadband Network (MSBN) as described in TR-134. The B reference point describes the interaction between the RADIUS clients in the vBG_MUX and vBG and the AAA server in the MSBN. Likewise, the R reference point describes the interaction between the vBG_MUX and vBG and the BPCF.

The following sections identify requirements that can be enabled by means of these reference points:

- Section 6.3.1.2 Connectivity monitoring by the vBG and vBG_MUX
- Section 6.4 Performance Monitoring requirements
- Section 6.5.2 QoS requirements on the vBG
- Section 6.9 AAA requirements

Detailed policy aspects and attributes is out of scope of the current TR-328 revision.

5.1.4 pBG Policy Reference Point (Rc)

The Rc reference point describes the interaction between the Policy Client in the pBG and the PDP.

The following sections identify requirements that can be enabled by means of this reference point:

- Section 6.2 Multi-homing
- Section 6.3.1.1 Connectivity monitoring by the pBG
- Section 6.3.2 LSL Failure and Protection
- Section 6.4 Performance Monitoring requirements
- Section 6.5.1 QoS requirements on the pBG
- Section 6.9 AAA requirements

Detailed policy aspects and attributes is out of scope of the current TR-328 revision.

 © The Broadband Forum. All rights reserved

5.1.5 SDN Client Reference Point (RCI)

The RCI reference point describes the interaction between the SDN Client in the pBG, vBG and vBG Hosting Infrastructure to the SDN Controller used to control resources in the pBG and vBG.

SDN control is out of scope of the current TR-328 revision.

5.1.6 Customer portal for Virtual Business Gateway

Today's physical Business Gateways include a management plane that enables customer administrators to monitor and configure their devices. When using the vBG, a similar level of control must be provided to the customer. Unlike the Business Gateway, this control plane is decoupled from the physical device. Hence, a customer portal is introduced in this Technical Report as a logical function that can support similar tasks to a Business Gateway management interface, such as:

- Self-Care: service selection, service configuration: addressing, routing, security, …
- Monitoring: availability, statistics, historical reports, logs, …

The customer portal provides additional features that are possible based on the virtual nature of vBG function, namely:

- Abstraction: the customer portal can monitor and configure multiple customer sites, instead of a specific site device.
- Simplification: the customer portal is not bound to a specific vendor device management interface, a simple user interface can be provided.
- Service activation/modification: the customer is not limited by the capability of the physical BG. As long as services and processing capacity are available in the vBG, a customer can add a new service on demand, e.g. "add per application reporting" or "add firewall service"; or modify service parameters, e.g. "increase bandwidth from 30Mbps to 50Mbps").

The customer portal is not defined as a component of the VBG System, but its presence drives the need for pBG and vBG to accept on demand changes, such as configuration and policies, and to expose externally relevant monitoring information.

The vBG function will need to have safeguard mechanisms to protect from misconfigurations in the customer portal, which could result in service impact for the customer or harm the SP's infrastructure.

Further specification of the portal capabilities and requirements is out of scope for this document.

 © The Broadband Forum. All rights reserved

5.2 Management of advanced forwarding use cases

For providing advanced services, the vBG/pBG have to apply different processing on specific flows:

- A flow may represent a customer organization (e.g. IP subnet), an employee or device (e.g. IP address), a type of application (e.g. a 5-tuple).
- Differentiated processing may include security, QoS, routing, chaining of value added services, statistics, etc.

The section below provides two example use cases that require processing traffic at flow level, involving coordination between the pBG and the vBG.

Use case 1: Blocking an Undesired Flow

Figure 23 depicts the case where some undesired traffic needs to be blocked at the pBG, for example a denial of service flow:

- Step 1: The undesired flow is detected by the IDS
- Step 2: The IDS function informs the controller that this flow needs to be blocked. The controller pushes a rule or a set of rules to block the flow on the pBG
- Step 3: Upon receipt, the pBG applies this rule to block the flow.

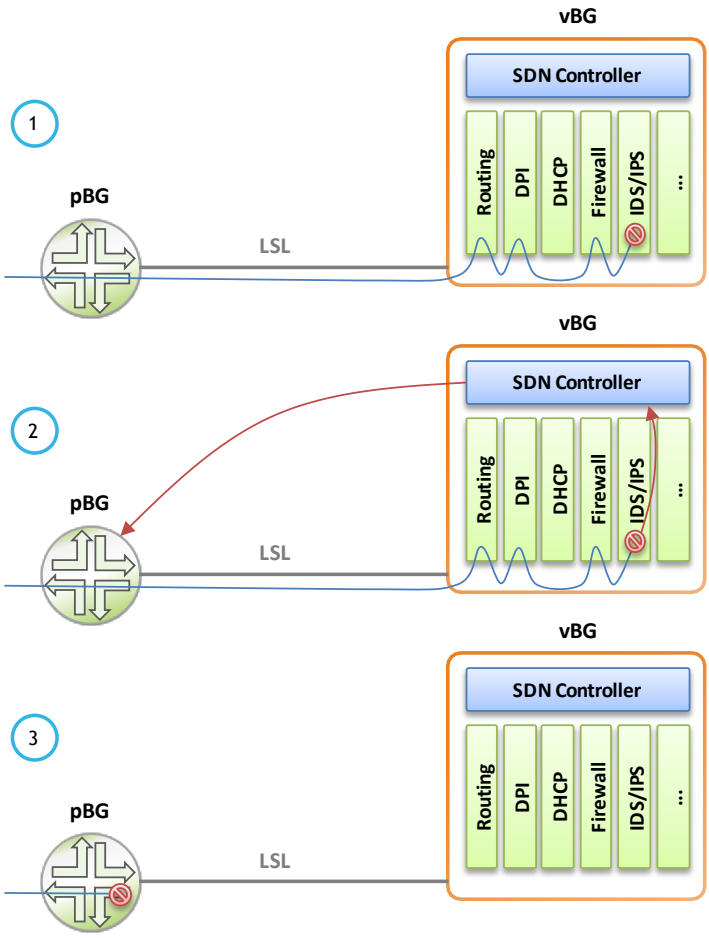

Figure 23 – Flow control example – Pushing a blocking rule in the pBG

© The Broadband Forum. All rights reserved

Use case 2: CoS change

Figure 24 depicts a use case describing a change of Class of Service on the pBG upon a decision taken by a DPI function hosted in the vBG:

- Step 1: A given flow is assigned to the default Best Effort Class of Service on the pBG-LSL Interface
- Step 2: The DPI function recognizes a top priority application and remaps the flow or the set of flows to a higher Class of Service in the vBG. The DPI function also informs the controller to reflect this change of CoS on the pBG. Hence, the controller pushes the appropriate rule or set of rules to change the CoS of the flow on the pBG.
- Step 3: The flow is now directly treated with the higher CoS on the pBG and forwarded accordingly.

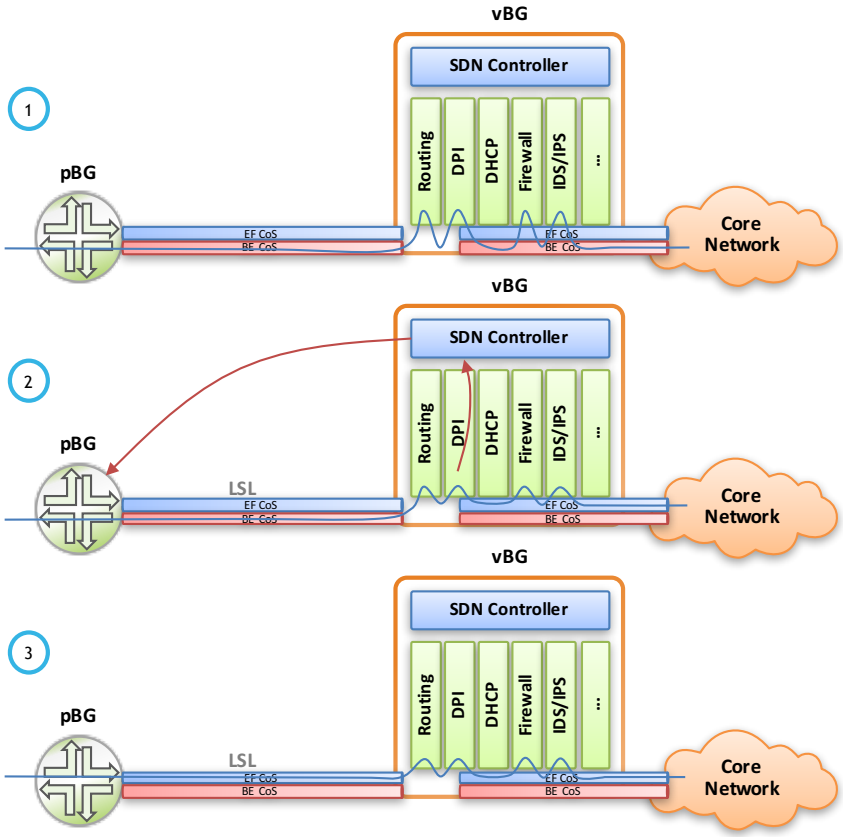

Figure 24 – Flow control example – Changing CoS of a flow in the pBG

In order to support these and other advanced use cases, the pBG must operate consistently with the vBG and as such a common control plane must be available. There are several options to fulfil this requirement, based on different systems:

- SDN control framework
- Policy domain
- Management systems

© The Broadband Forum. All rights reserved

In all cases, it allows pBG and vBG to report events that can then be used as triggers for forwarding changes in the nodes, enforced by the SDN control framework, the policy domain or the network element managers respectively.

In any case, a workflow will be required to respond to trigger events with a given action or set of actions. If different platforms are receiving the trigger events than those enforcing the actions, then a higher degree of integration will be needed.

For the specific cases described above, the following is a potential high level workflow. Upon detection of the trigger event at the vBG:

- **SDN control framework:** the vBG reports the new status (e.g. a given flow needs to be blocked or assigned a different CoS) to the SDN control framework. Based on this, the SDN control framework takes action and installs or modifies the forwarding rule for that given flow in the pBG.
- **Policy domain:** the vBG reports the event to a northbound application, that may in turn use the policy domain to distribute a new forwarding behavior for a given flow.
- **Management systems:** the vBG reports an alarm to the management system. The management system processes the alarm and takes action based on its content, proceeding to program the pBG to execute the expected action, e.g. using TR-069 [1], NETCONF, SNMP, etc.

 © The Broadband Forum. All rights reserved

6 VBG System requirements

6.1 End-to-end network requirements

As described in section 4.3.4, the LSL that connects a pBG to its vBG is composed of two segments: pBG to vBG_MUX, and vBG_MUX to vBG. The LSL segment from the pBG to vBG_MUX can either be:

- Flat LSL: a native Ethernet connection, statically or dynamically provisioned through the MSBN, as described in section 6.1.1.
- Overlay LSL: an Ethernet connection, statically or dynamically established over an IP network using tunneling techniques, as described in section 6.1.2.

The Flat Ethernet LSL scenario is convenient for Service Providers who already have a 1:1 VLAN access and backhaul architecture in place. It may also allow reusing deployed Business 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. In the case of Overlay LSL, the LSL segment between the vBG_MUX and the vBG still requires maintaining segregation between the different enterprises and branches. The choice of network technology (e.g. VLAN, VXLAN, etc.) for the segment between the vBG_MUX and the vBG is left up to implementation and is out of scope of TR-328.

6.1.1 Flat LSL connectivity

In the flat model, the pBG is connected to the vBG hosting infrastructure either directly or through an Ethernet access/aggregation network using the 1:1 VLAN model, documented in TR-101 [2] and TR-178 [6].

In addition to the direct Ethernet connection model, the pBG may be connected to the vBG hosting infrastructure through an existing MS-BNG. In that case, the MS-BNG may be statically or dynamically provisioned to provide a L2 extension between the pBG access line and the vBG_MUX, as described in section 7.1.1/TR-317. AAA related requirements are described in section 6.9.

6.1.2 Overlay LSL connectivity

TR-317 describes the use of GRE tunneling for Overlay LSL. While GRE tunneling will work for the vast majority of VBG System deployments, it does not provide NAT traversal capabilities. For GRE to work through NAT, a 1:1 mapping should be provided, which defeats the purpose of most NAT deployments (using NAPT).

 © The Broadband Forum. All rights reserved

The following are a couple of examples of network environments where NAT traversal will be essential:

- pBG WAN connection is terminated in a MS-BNG that performs NAPT
- pBG providing VBG System services is deployed behind the existing CPE for cost / time to market reasons

In order to successfully traverse a network where NAPT is deployed between the pBG and the vBG, a tunneling technology with a transport layer (e.g. UDP) is required. In addition to Ethernet over GRE, this Technical Report describes two alternatives to support the case when the LSL must cross a NAT function:

- Ethernet over L2TPv3 over UDP
- Ethernet over VXLAN

Choice of one or the other will depend on pBG capabilities and Service Provider preference.

Note that in addition to providing NAT traversal capabilities, the use of tunneling technologies over a transport layer protocol can also be used to improve traffic load balancing in the network. As an example, this can be done by hashing the headers of the customer traffic and setting the result as the source UDP port of the external encapsulation, which results in increased entropy for hashing by the network elements in the path between the pBG and the vBG.

[R-1] The pBG MUST support at least one of the following Overlay LSL technologies:
- Ethernet over GRE
- Ethernet over L2TPv3 over UDP
- Ethernet over VXLAN

[R-2] The vBG_MUX MUST support all the following Overlay LSL technologies:
- Ethernet over GRE
- Ethernet over L2TPv3 over UDP
- Ethernet over VXLAN

[R-3] The vBG_MUX SHOULD support prefix-list filtering of allowed LSL tunnels sources.

 © The Broadband Forum. All rights reserved

6.1.2.1 Overlay LSL using GRE

GRE is described in RFC 2784 [24], and provides a generic format for encapsulating and transporting payloads of multiple protocols, including Ethernet. GRE is normally transported over IPv4 or IPv6 natively, without transport (i.e. UDP/TCP) header. Hence, it does not provide NAT traversal capabilities.

TR-328 makes use of stateless GRE tunnels. This allows minimizing resource consumption at the pBG and the vBG_MUX, which allows more cost effective pBG models and more scalable vBG_MUX implementations.

The provisioning of the LSL settings in the pBG will be done using one of the following mechanisms:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Provisioned by the ACS, using TR-069
- Obtained from the network via DHCP

Provisioning of GRE tunnels via TR-069 is supported using the TR-181i2 [7] data model.

The provisioning of the LSL settings in the vBG_MUX will use one of the following methods:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Dynamically learned from the received GRE encapsulated packets from the pBG and then used for downstream traffic to the pBG

Figure 25 shows the resulting high level encapsulation in the LSL when using this tunneling method.

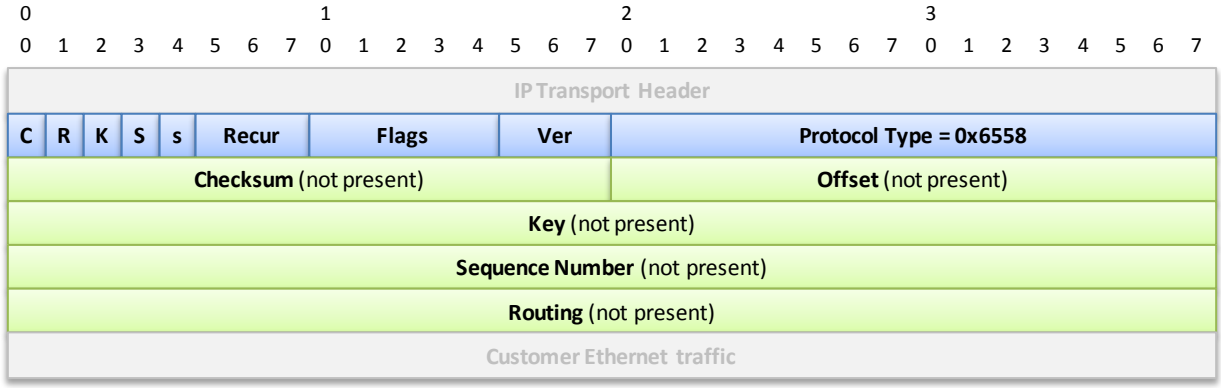

Figure 25 – LSL encapsulation for Ethernet over GRE

 © The Broadband Forum. All rights reserved

Table 2 describes the values that should be set in each of the headers of the GRE encapsulation in the LSL.

Header Field	Value
Source IP address	**pBG to vBG:** pBG WAN IP **vBG to pBG:** vBG_MUX IP
Destination IP address	**pBG to vBG:** vBG_MUX IP **vBG to pBG:** pBG WAN IP
IP Protocol Type / Next-Header	GRE (0x2F)
GRE Protocol Type	Transparent Ethernet Bridging (0x6558)
C/R/K/S/s bits	Set to 0
Recursion Control	Set to 0
Flags	Set to 0
Version	0
Offset/Key/Sequence Number/Routing	Not present
Source MAC address	**pBG to vBG:** End-user device MAC **vBG to pBG:** vBG MAC address *
Destination MAC address	**pBG to vBG:** vBG MAC address * **vBG to pBG:** End-user device MAC In both cases, the broadcast MAC is used for broadcast traffic, e.g. for ARPs

Table 2 – LSL settings for GRE

Note: the vBG MAC address in Table 2 refers to the MAC address of the target vBG system component in the case of a distributed vBG (e.g. DHCP Server, Firewall, etc.)).

The following requirements apply to the pBG when using Ethernet over GRE tunnels for the LSL:

[R-4] The pBG MUST support stateless GRE tunnels (no signaling required)
[R-5] The pBG MUST support stateless GRE tunnels using IPv4 encapsulation.
[R-6] The pBG SHOULD support stateless GRE tunnels using IPv6 encapsulation.
[R-7] The pBG MUST support bridging Ethernet frames into a GRE tunnel.
[R-8] The pBG MUST support using the LSL settings in Table 2.
[R-9] The pBG MUST support static provisioning of GRE LSL settings
[R-10] The pBG SHOULD support TR-069 provisioning of GRE LSL settings
[R-11] The pBG SHOULD support obtaining GRE LSL settings via DHCP, as described in section 6.1.4.
[R-12] Upon receiving downstream encapsulated traffic from the vBG, the pBG MUST:
- Decapsulate GRE
- If the Protocol Type in GRE header is Transparent Ethernet Bridging (0x6558), then it must process the 802.3 [13] frame following the GRE header.
- The frame should be forwarded per the MAC forwarding table.

 © The Broadband Forum. All rights reserved

The following requirements apply to the vBG_MUX when using GRE tunnels for the LSL:

[R-13] The vBG_MUX MUST support stateless GRE tunnels (no signaling required)
[R-14] The vBG_MUX MUST support stateless GRE tunnels using IPv4 encapsulation.
[R-15] The vBG_MUX SHOULD support stateless GRE tunnels using IPv6 encapsulation.
[R-16] The vBG_MUX MUST support bridging Ethernet frames into a GRE tunnel.
[R-17] The vBG_MUX MUST support using the LSL settings in Table 2.
[R-18] The vBG_MUX MUST support static provisioning of GRE LSL settings.
[R-19] The vBG_MUX SHOULD support dynamically learning the GRE LSL settings from encapsulated packets received from the pBG. Learned encapsulation is then used on downstream traffic to the pBG.
[R-20] Upon receiving upstream encapsulated traffic from the pBG, the vBG_MUX MUST:

- Decapsulate GRE
- If the Protocol Type in GRE header is Transparent Ethernet Bridging (0x6558), then it must process the 802.3 frame following the GRE header.
- The frame should be forwarded to the selected vBG for this pBG.

 © The Broadband Forum. All rights reserved

6.1.2.2 Overlay LSL using L2TPv3 over UDP

L2TPv3 is described in RFC 3931 [30] and allows its transport over IP or UDP. L2TPv3 is expanded in RFC 4719 [32] to allow for the transport of L2 frames across over L2TPv3. Since the main objective of supporting additional tunneling technologies besides GRE in TR-328 is to be able to support NAT traversal, TR-328 supports L2TPv3 over UDP tunnels for the LSL, for the transport of Ethernet customer frames between the pBG and the vBG.

While the L2TPv3 specification includes both control plane and data plane, TR-328 will make use of stateless L2TPv3 tunnels. This allows minimizing resource consumption at the pBG and the vBG_MUX, which allows more cost effective pBG models and more scalable vBG_MUX implementations.

The provisioning of the LSL settings in the pBG will be done using one of the following mechanisms:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Provisioned by the ACS, using TR-069
- Obtained from the network via DHCP

At the time of this writing, work is ongoing on a TR-181i2 data model for TR-069 provisioning of L2TPv3 over UDP tunnels. Meanwhile, vendor-specific extensions can be used.

The provisioning of the LSL settings in the vBG_MUX will use one of the following methods:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Dynamically learned from the received L2TPv3 over UDP encapsulated packets from the pBG and then used for downstream traffic to the pBG

Figure 25 shows the resulting high level encapsulation in the LSL when using this tunneling method.

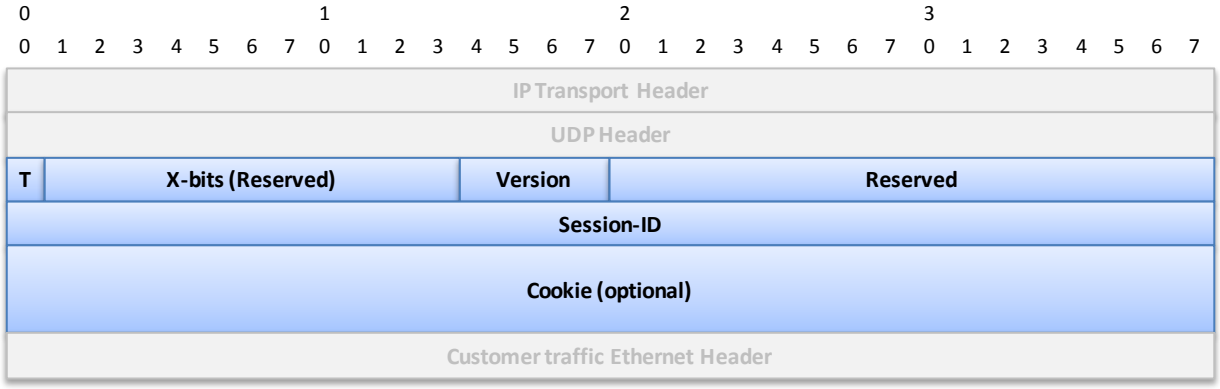

Figure 26 – LSL encapsulation for Ethernet over L2TPv3oUDP

 © The Broadband Forum. All rights reserved

Table 3 describes the values that should be set in each of the headers of the L2TPv3oUDP encapsulation in the LSL.

Header Field	Value
Source IP address	**pBG to vBG:** pBG WAN IP **vBG to pBG:** vBG_MUX IP
Destination IP address	**pBG to vBG:** vBG_MUX IP **vBG to pBG:** pBG WAN IP
IP Protocol Type / Next-Header	UDP (0x11)
Source UDP Port	Configurable (1701 recommended)
Destination UDP Port	1701
T-bit	Set to 0 (data message)
X Bits	Set to 0
Reserved field	Set to 0
Session ID	Configurable. Default=1.
Cookie	Configurable.
Source MAC address	**pBG to vBG:** End-user device MAC **vBG to pBG:** vBG MAC address *
Destination MAC address	**pBG to vBG:** vBG MAC address * **vBG to pBG:** End-user device MAC In both cases, the broadcast MAC is used for broadcast traffic, e.g. for ARPs

Table 3 – LSL settings for L2TPv3oUDP

Note: the vBG MAC address in Table 3 refers to the MAC address of the target vBG system component in the case of a distributed vBG (e.g. DHCP Server, Firewall, etc.)).

The following requirements apply to the pBG when using L2oL2TPv3 tunnels for the LSL:

[R-21] The pBG MUST support stateless L2TPv3 over UDP tunnels (no signaling required)
[R-22] The pBG MUST support stateless L2TPv3 over UDP tunnels using IPv4 encapsulation.
[R-23] The pBG SHOULD support stateless L2TPv3 over UDP tunnels using IPv6 encapsulation.
[R-24] The pBG MUST support bridging Ethernet frames into a L2TPv3 tunnel (L2oL2TPv3).
[R-25] The pBG MUST support using the LSL settings in Table 3.
[R-26] The pBG MUST support setting the pBG MAC address in the six lower significant bytes of the Cookie field.
[R-27] The pBG MUST support static provisioning of L2TPv3 over UDP LSL settings
[R-28] The pBG SHOULD support TR-069 provisioning of L2TPv3 over UDP LSL settings
[R-29] The pBG SHOULD support obtaining L2TPv3 over UDP LSL settings via DHCP, as described in section 6.1.4.
[R-30] Upon receiving downstream encapsulated traffic from the vBG, the pBG MUST:
- Decapsulate L2TPv3
- If the Protocol Type in IP header is UDP (0x11) and the UDP Destination Port is 1701, then it must process the 802.3 frame following the L2TPv3 header.
- The frame should be forwarded per the MAC forwarding table.

 © The Broadband Forum. All rights reserved

The following requirements apply to the vBG_MUX when using L2oL2TPv3 tunnels for the LSL:

[R-31] The vBG_MUX MUST support stateless L2TPv3 over UDP tunnels (no signaling required)
[R-32] The vBG_MUX MUST support stateless L2TPv3 over UDP tunnels using IPv4 encapsulation.
[R-33] The vBG_MUX SHOULD support stateless L2TPv3 over UDP tunnels using IPv6 encapsulation.
[R-34] The vBG_MUX MUST support bridging Ethernet frames into a L2TPv3 tunnel (L2oL2TPv3).
[R-35] The vBG_MUX MUST support using the LSL settings in Table 3.
[R-36] The vBG_MUX MUST support static provisioning of L2TPv3 over UDP LSL settings.
[R-37] The vBG_MUX SHOULD support dynamically learning the L2TPv3 over UDP LSL settings from encapsulated packets received from the pBG. Learned encapsulation is then used on downstream traffic to the pBG.
[R-38] Upon receiving upstream encapsulated traffic from the pBG, the vBG_MUX MUST:

- Decapsulate L2TPv3
- If the Protocol Type in IP header is UDP (0x11) and the UDP Destination Port is 1701, then it must process the 802.3 frame following the L2TPv3 header.
- The frame should be forwarded to the selected vBG for this pBG.

 © The Broadband Forum. All rights reserved

6.1.2.3 Overlay LSL using VXLAN

VXLAN is described in RFC 7348 [38]. It allows the transport of Ethernet frames and it is encapsulated in UDP, allowing for NAT traversal.

Figure 27 depicts the use of VXLAN for Overlay LSL connectivity.

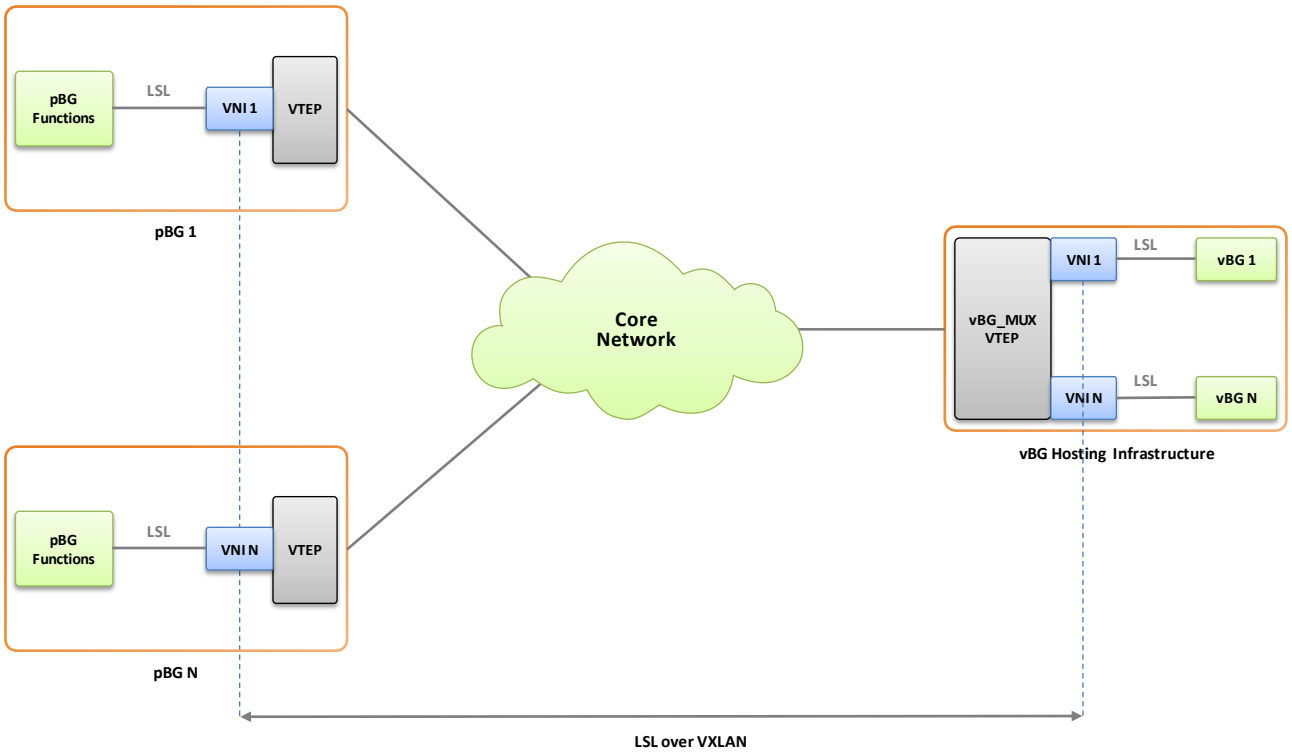

Figure 27 – LSL connectivity using VXLAN

As shown in Figure 27, the pBG and vBG hosting infrastructure are VXLAN Tunnel End Points (VTEPs), performing the encapsulation/decapsulation of customer traffic in VXLAN. In this architecture, a unique VXLAN Network Identifier (VNI) is assigned per vBG by the service provider and used in the VXLAN LSL encapsulation between the pBG and the vBG_MUX.

While using a control plane for VXLAN (e.g. EVPN) is possible and even highly desirable in some environments, TR-328 focusses on point-to-point LSL tunnels, where the tunnel endpoints are clear and well defined, the pBG and the vBG_MUX.

As in the case of GRE and L2TPv3 over UDP, the provisioning of the LSL settings for VXLAN in the pBG will be done using one of the following mechanisms:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Provisioned by the ACS, using TR-069
- Obtained from the network via DHCP

At the time of this writing, work is ongoing on a TR-181i2 data model for TR-069 provisioning of VXLAN tunnels in the pBG. Meanwhile, vendor-specific extensions can be used.

© The Broadband Forum. All rights reserved

The provisioning of the LSL settings for VXLAN in the vBG_MUX will use one of the following methods:

- Provisioned by the EMS/NMS, using NETCONF/YANG or SNMP
- Dynamically learned from the received VXLAN encapsulated packets from the pBG and then used for downstream traffic to the pBG

Table 4 describes the values that should be set in each of the headers of the VXLAN encapsulation in the LSL.

Header Field	Value
Source IP address	**pBG to vBG:** pBG WAN IP **vBG to pBG:** vBG_MUX IP
Destination IP address	**pBG to vBG:** vBG_MUX IP **vBG to pBG:** pBG WAN IP
IP Protocol Type / Next-Header	UDP (0x11)
Source UDP Port	Configurable (4789 recommended)
Destination UDP Port	4789
VXLAN Network Identifier	Configurable (one per enterprise)
Source MAC address	**pBG to vBG:** End-user device MAC **vBG to pBG:** vBG MAC address *
Destination MAC address	**pBG to vBG:** vBG MAC address * **vBG to pBG:** End-user device MAC In both cases, the broadcast MAC is used for broadcast traffic, e.g. for ARPs

Table 4 – LSL settings for VXLAN

Note: the vBG MAC address in Table 4 refers to the MAC address of the target vBG system component in the case of a distributed vBG (e.g. DHCP Server, Firewall, etc.)).

Special care should be taken when migrating a vBG from one host to another, resulting of VM relocation. This could be within a PoP or data center or to a different PoP or data center. Reprogramming of the pBG to change the destination VTEP of the new vBG_MUX can be done in several ways. vBG mobility is outside of the scope of this TR-328 issue. The following are a couple of examples:

- By means of orchestration: when mobility of a vBG is detected at the orchestration layer, the orchestrator triggers reconfiguration of the destination tunnel endpoint at the pBG, by means of TR-069, NETCONF, etc.
- By means of SDN Control: when the vBG comes up in a new host, the SDN controller programs the destination tunnel endpoint at the pBG to the new host.

 © The Broadband Forum. All rights reserved

The following requirements apply to the pBG when using VXLAN tunnels for the LSL:

[R-39] The pBG MUST support VXLAN tunnels
[R-40] The pBG MUST support VXLAN tunnels using IPv4 encapsulation.
[R-41] The pBG SHOULD support VXLAN tunnels using IPv6 encapsulation.
[R-42] The pBG MUST support bridging Ethernet frames into a VXLAN tunnel.
[R-43] The pBG MUST support using the LSL settings in Table 4.
[R-44] The pBG MUST support static provisioning of VXLAN LSL settings
[R-45] The pBG SHOULD support TR-069 provisioning of VXLAN LSL settings
[R-46] The pBG SHOULD support obtaining VXLAN LSL settings via DHCP, as described in section 6.1.4.
[R-47] Upon receiving downstream encapsulated traffic from the vBG, the pBG MUST:
- Decapsulate VXLAN
- If the Protocol Type in IP header is UDP (0x11) and the UDP Destination Port is 4789, then it must process the 802.3 frame following the VXLAN header.
- The frame should be forwarded per the MAC forwarding table, if matching the VNI configured for the LSL.

The following requirements apply to the vBG_MUX when using VXLAN tunnels for the LSL:

[R-48] The vBG_MUX MUST support stateless VXLAN tunnels
[R-49] The vBG_MUX MUST support stateless VXLAN tunnels using IPv4 encapsulation.
[R-50] The vBG_MUX SHOULD support stateless VXLAN tunnels using IPv6 encapsulation.
[R-51] The vBG_MUX MUST support bridging Ethernet frames into a VXLAN tunnel.
[R-52] The vBG_MUX MUST support using the LSL settings in Table 4.
[R-53] The vBG_MUX MUST support static provisioning of VXLAN settings.
[R-54] The vBG_MUX SHOULD support dynamically learning the VXLAN LSL settings from encapsulated packets received from the pBG. Learned encapsulation is then used on downstream traffic to the pBG.
[R-55] Upon receiving upstream encapsulated traffic from the pBG, the vBG_MUX MUST:
- Decapsulate VXLAN
- If the Protocol Type in IP header is UDP (0x11) and the UDP Destination Port is 4789, then it must process the 802.3 frame following the VXLAN header.
- The frame should be forwarded to the selected vBG for this pBG, based on the VNI.

 © The Broadband Forum. All rights reserved

6.1.3 Multi-VLAN LSL

Transporting multiple VLANs on the LSL interface will be required for several use cases:

- Multiple LAN segments in a bridged pBG
- Multiple services (e.g. routing contexts/VRFs) in a routed pBG
- Separate management interface
- Co-existence with legacy services, where VBG System services are provided over a dedicated VLAN in the WAN interface, while keeping existing services running

In the case of Flat LSL, the pBG will forward the traffic with the corresponding VLAN tags over the pBG-LSL interface. In the case of Overlay LSL, the pBG will use a single tunnel and transport the required VLANs over the tunnel header.

6.1.3.1 pBG Multi-VLAN LSL requirements

The following general requirements apply to the pBG, in all forwarding modes:

[R-56] The pBG MUST support an 802.1Q [11] VLAN tagged pBG-LSL interface
[R-57] The pBG MUST support multiple VLANs on the pBG-LSL interface

In the case of Overlay LSL, the following additional requirement applies to the pBG:

[R-58] The pBG MUST support multiple VLANs over an Overlay LSL tunnel

The following requirements apply to a Bridged pBG:

[R-59] The pBG MUST support bridging multiple VLANs
[R-60] The pBG MUST support configuring a list of allowed VLANs on a LAN port
[R-61] The pBG MUST support configuring the default VLAN identifier (VID) of a LAN port, for untagged traffic
[R-62] The pBG MUST support assigning a Wi-Fi SSIDs to a VLAN

The following requirements apply to a Routed pBG:

[R-63] The pBG MUST support multiple routing contexts
[R-64] The pBG MUST support configuring the VID to use in the pBG-LSL interface, for a given routing context
[R-65] The pBG MUST support multiple IP interfaces in a routing context
[R-66] The pBG MUST support assigning a LAN port to an IP interface
[R-67] The pBG MUST support assigning a sub-interface of a LAN port to an IP interface
[R-68] The pBG MUST support assigning a Wi-Fi SSID to an IP interface

The following requirements apply to a Routing & Bridging pBG:

[R-69] The pBG MUST support requirements [R-59] to [R-62]
[R-70] The pBG MUST support requirements [R-63] to [R-65]

 © The Broadband Forum. All rights reserved

[R-71] The pBG MUST support assigning a bridged VLAN to an IP interface

6.1.3.2 vBG and vBG_MUX Multi-VLAN LSL requirements

The following requirements apply to the vBG:

[R-72] The vBG MUST support an 802.1Q VLAN tagged vBG-LSL interface
[R-73] The vBG MUST support multiple VLANs on the vBG-LSL interface
[R-74] The vBG MUST support multiple routing contexts
[R-75] The pBG MUST support assigning one or more VIDs to a routing context, one IP interface per VID

In the case of Overlay LSL, the following requirements apply to the vBG_MUX:

[R-76] The vBG_MUX MUST support multiple VLANs over an Overlay LSL tunnel
[R-77] The vBG_MUX MUST support transparently forwarding the VLAN-encapsulated traffic between the pBG and the vBG

6.1.4 pBG Overlay LSL tunnel attributes via DHCP

The pBG may use DHCP to obtain its IP address(es) used for the device, either as a host in the LAN when the pBG operates in bridged mode, or for its WAN interface, when the pBG operates in routed or routed & bridged mode. TR-124 provides a set of requirements for the functionality provided by the pBG as described in the section titled WAN.DHCPC.

To setup the Overlay LSL, the pBG must be provided with the necessary tunneling information. In the case the information is provided via DHCP, Table 5 lists the required DHCPv4 and DHCPv6 options to be used for this purpose.

Tunnel Attribute	DHCPv4 Option	DHCPv6 Option	DHCP sub-option
Tunnel-Type	125	17	21
Server-Endpoint	125	17	22
Client-Endpoint	125	17	23
Tunnel-Specific	125	17	20

Table 5 – Overlay LSL DHCPv4 and DHCPv6 options

 © The Broadband Forum. All rights reserved

Table 6 describes the values that should be set for each of the Overlay LSL tunnel attributes, when passing that information to the pBG using DHCP options.

Tunnel Attribute	Value
Tunnel-Type	0: GRE 1: VXLAN 2: L2TPv3oUDP
Server-Endpoint	vBG_MUX tunnel endpoint IP address
Client-Endpoint	pBG tunnel endpoint IP address
Tunnel-Specific	String containing values for the configurable attributes of the tunnel type in use: • GRE: Not present • VXLAN: Contains the source UDP port and the VXLAN VNI. Format: "sport=<source-udp-port>;vni=<vni-value>" • L2TPv3oUDP: Contains the source UDP port, the Session ID and the Cookie Format: "sport=<source-udp-port>;sid=<session-id>;cookie=<cookie-value>"

Table 6 – Overlay LSL tunnel attribute values

The pBG must meet the following set of requirements:

[R-78] The pBG MUST support a DHCP Client that complies with the TR-124 WAN.DHCPC requirements.
[R-79] The pBG MUST be able to obtain the IP configuration of its network interface, through DHCP, prior to tunnel establishment.
[R-80] The pBG MUST be able to establish an Overlay LSL tunnel over IP to the vBG_MUX using the information received via DHCP, using the DHCP options described in Table 5 and the values listed in Table 6.
[R-81] The pBG's DHCP client MUST insert the Device-Type information in the dedicated Option 125 sub-option 24 in its all its DHCP messages
[R-82] The pBG MUST set the default value of DHCP option Device-Type MUST be the string "pBG" (without the quotes)
[R-83] The value of DHCP option Device-Type set by the pBG MUST be configurable.

6.1.5 MTU considerations

When using Overlay LSL, the network nodes in access and backhaul networks need to support a Maximum Transmission Unit (MTU) large enough to allow the transport of Ethernet over the overlay LSL without causing fragmentation.
[R-84] The MTU at the pBG MUST be configurable.
[R-85] The pBG MUST support an MTU that permits Ethernet frames of at least 1518 bytes to be encapsulated within the LSL tunnel.
[R-86] The MTU at the vBG_MUX MUST be configurable.
[R-87] The vBG_MUX MUST support an MTU that permits Ethernet frames of at least 1518 bytes to be encapsulated within the LSL tunnel.

 © The Broadband Forum. All rights reserved

6.1.6 Multiple pBG support

For redundancy reasons, as well a supporting geographically dispersed VPNs sites such as different customer sites or just separate buildings on the same campus, the vBG needs to support multiple pBG's.

[R-88] A single vBG MUST be able to support multiple pBGs, which belong to the same enterprise.

6.2 Multi-homing

Reliability is one of the main concerns for business customers. In order to ensure that they have a reliable connection to the Internet and/or to other business VPN sites, the business branch can have multiple uplink connections. Instead of being connected to the network via just one uplink, the pBG has two or more access links and can be connected to the vBG via two or more network accesses via one or more network providers. To increase the reliability of the connection to the vBG, those connections are typically running on separate physical mediums, that is, separate cables, fibers, radio interfaces or any combination of them.

In the current revision of TR-328 one of the access links will be considered primary and any additional access links will be used for backup purposes, providing higher reliability.

While vBG redundancy/high availability is possible and even desirable, it is left to implementation and it is out of scope of this document.

If the pBG supports multi-homing, the following general requirements apply:
[R-89] The pBG MUST have two or more WAN interfaces.
[R-90] The pBG MUST support setting an LSL to the vBG over any of its WAN interfaces.

6.2.1 Multi-homing to a single network

When the two or more access links of the pBG are multi-homed to the same network, several aspects are addressed implicitly:

- LSL authentication: The access links belongs to the same service provider that is providing the VBG System service to the customer. Implicit authentication of the pBG to vBG connection can be assumed, based on the authentication of the links of the pBG to the access network, e.g. MS-BNG authentication.
- pBG to vBG IP connectivity: The service provider may choose either public or private IP addressing for the pBG access links and the vBG, as the packets going from the pBG to the vBG will not traverse a third-party network.

 © The Broadband Forum. All rights reserved

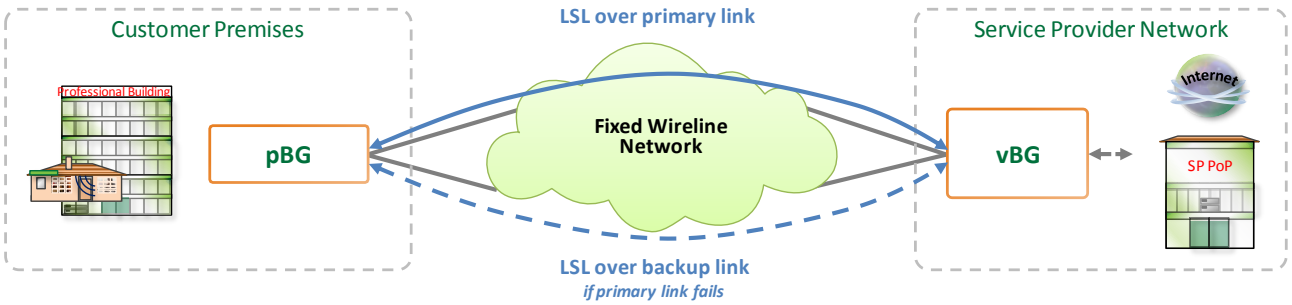

Figure 28 – Multi-homing to a single network

6.2.2 Multi-homing to multiple networks of the same service provider

This case is very similar to multi-homing to a single network. Care should be taken if both networks are using RFC 1918 [16] addressing, as issues may arise if they have partial or total overlapping of IP addressing.

An additional consideration is that it may be complex, or not even possible, to make use of Flat LSL connectivity between the pBG and the vBG. A good example of such case is the use of a 3GPP network to provide a backup access, as a 3GPP network normally provides a L3 connection terminated at the Packet Data Network Gateway (PGW). In such cases, the use of Overlay LSL will be required.

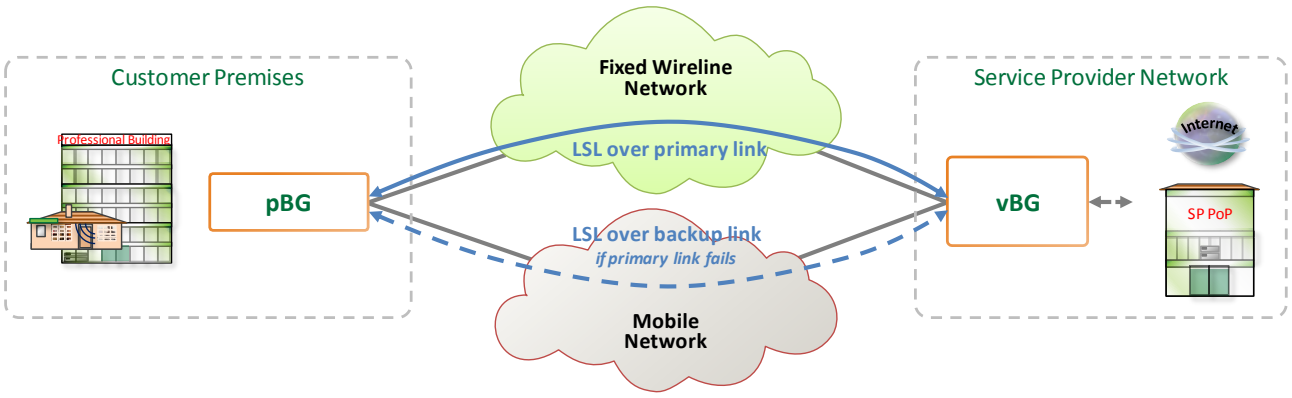

Figure 29 – Multi-homing to multiple networks of the same service provider

6.2.3 Multihoming to multiple networks of different service providers

This describes the case where at least one of the pBG access links is over a network service provider different than the VBG System service provider.

Even though it would be possible to get a Layer 2 service connecting the pBG to the vBG from a network service provider other than the VBG System service provider, so that a Flat LSL can be used, it will normally be easier and more cost-effective to use Overlay LSL.

In such case, it will be required to use public IP addressing for the vBG_MUX, so that it is reachable from a pBG connected to a third-party network. The pBG will also require a public IP

 © The Broadband Forum. All rights reserved

address, or using an LSL tunneling technology that allows for NAT traversal. See section 6.1.2 for details.

If basic authentication is required, the pBG could authenticate itself to a AAA server of the VBG System service provider using Extensible Authentication Protocl (EAP). The vBG_MUX could act as a RADIUS proxy, caching the authentication state of the pBG and its source IP address.

If stronger authentication and encryption are desired, IPsec could be used.

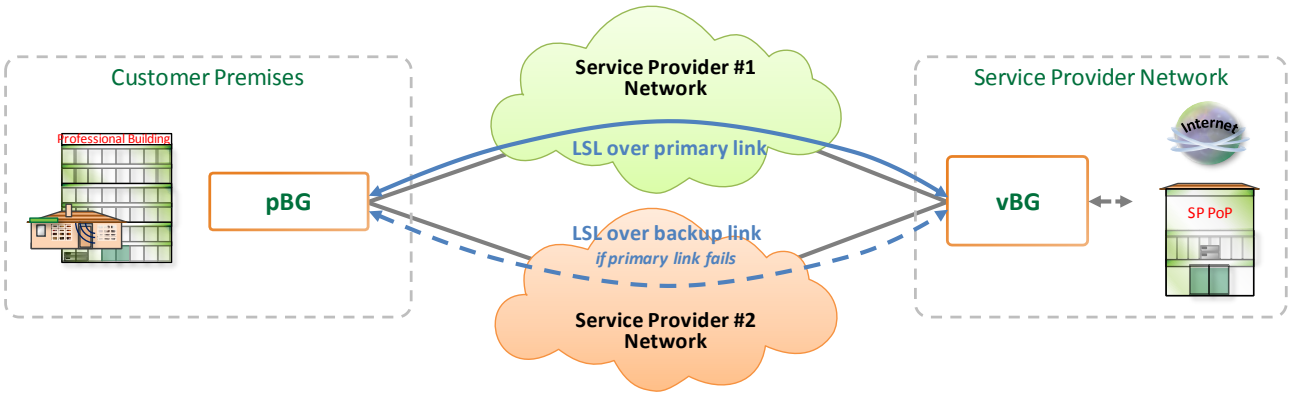

Figure 30 – Multi-homing to multiple networks of different service providers

6.3 LSL monitoring and protection

By virtualizing the BG and functionally splitting it in two components, the pBG and vBG, together with the shift of vital functions from the customer site to the network, failure of the LSL will result in partial or total failure of services delivered to and from the branch. Additional complexity is introduced by the fact that vBG may be distributed among multiple network elements, including equipment at either side of the LSL.

The following services could be impacted in case of LSL failure:

- Connectivity to the Internet
- Connectivity to other branches of the business customer
- Access to applications and services which are hosted in the vBG and/or the cloud
- Intra-office LAN communication issues, which could prevent:
  - Communication between end user devices in the branch
  - Printing documents on a network printer
  - etc.

The intra-office LAN issues could be caused by failure to resolve a machine/printer name due to lack of DNS services or by a device not being able to obtain an IP address due to lack of DHCP services, in the case of a bridged pBG.

Some of the issues listed above represent a regression compared to the present model of operation where, in case of a WAN failure, intra-office LAN connectivity is preserved. In addition, while

 © The Broadband Forum. All rights reserved

temporary outages of Internet access and cloud-hosted applications may be acceptable for residential subscribers, it may cause serious issues for a number of business customers.

6.3.1 Connectivity management and LSL monitoring

Both ends of the LSL need to monitor the availability of the connectivity between the pBG to the vBG. The connectivity monitoring is based on up to two layers: pBG to vBG (LSL scope) and in the case of Overlay LSL, pBG to vBG-MUX (Tunnel scope). The monitoring mechanism between the pBG and the vBG is based on ARP requests as described in TR-146 [5] for IP session monitoring. The tunnel monitoring is based on ICMP.

In cases of high availability there may be significance to the failure detection and service restoration time that are part of the Service Level Agreement (SLA). In such cases, BFD can be used instead of ARP for fast failure detection of the LSL. BFD asynchronous mode will be used; hence a BFD session is first set up between the pBG and the vBG over the LSL and then periodic hello messages will be sent by both nodes to keep the session alive.

The pBG uses the connectivity information to control the operation of the backup DHCP server in case where the DHCP server is not local, as per Section 6.1.3.2.1/TR-317. In cases where there is redundancy on the LSL the pBG uses this information for switching to the other LSL connection. In addition, the vBG may use this information for resource management.

Figure 31 below shows the ARP ping keep-alive mechanism on both pBG and vBG.

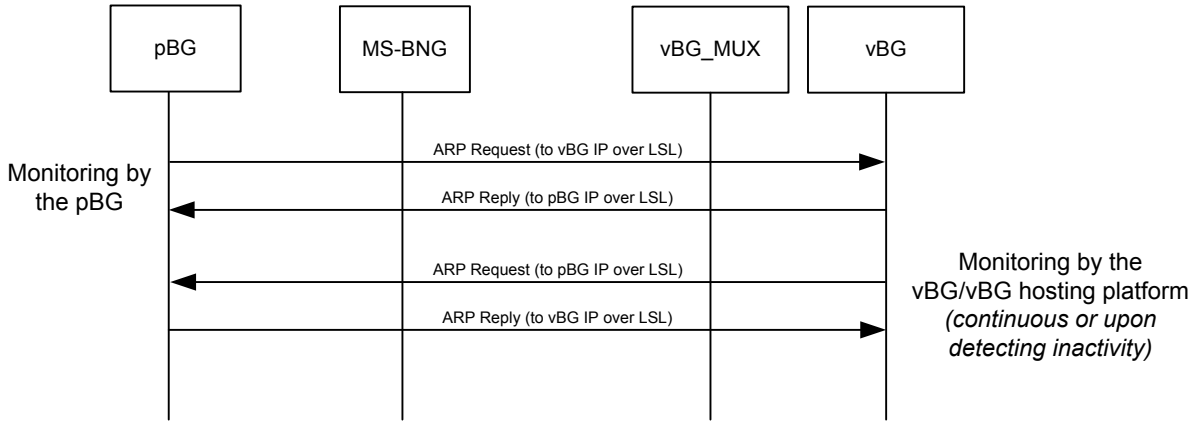

Figure 31 – LSL monitoring using ARP

 © The Broadband Forum. All rights reserved

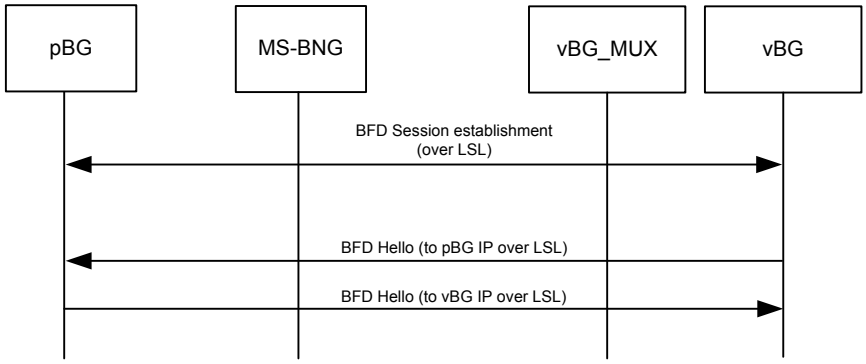

Figure 32 – LSL monitoring using BFD

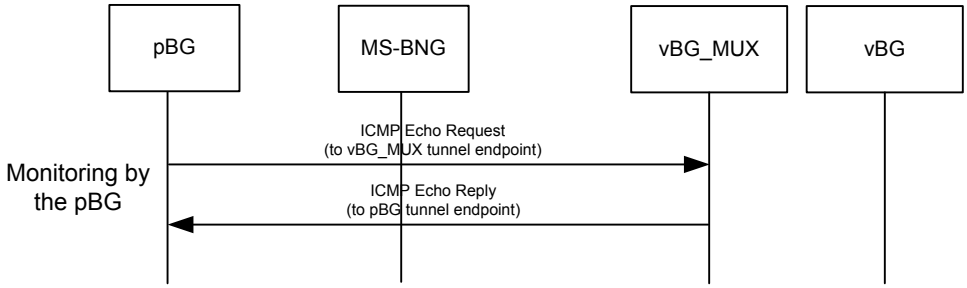

Figure 33 – pBG tunnel monitoring using ICMP

The following table lists the IP addresses to use in the pBG, vBG and vBG_MUX for LSL and tunnel monitoring, depending on the pBG forwarding mode:

Monitoring mechanism	LSL Type	pBG forwarding mode	pBG address	vBG_MUX address	vBG address
LSL Monitoring	Flat	Bridged	LAN IP	N/A	vBG-LSL IP *Gateway IP*
		Routed/IRB	pBG-LSL IP *WAN IP*		
	Overlay	Bridged	LAN IP		
		Routed/IRB	pBG-LSL IP		
Tunnel Monitoring	Overlay	Bridged	Underlay IP *WAN IP*	vBG_MUX tunnel endpoint	N/A
		Routed/IRB			

Table 7 – LSL and tunnel endpoint monitoring IP addressing

© The Broadband Forum. All rights reserved

6.3.1.1 Connectivity monitoring by the pBG

6.3.1.1.1 LSL monitoring requirements

The pBG monitors the state of its connectivity to the vBG via the LSL monitoring mechanism, based on ARP and/or Neighbor Discovery (ND). BFD can be used instead of ARP/ND for fast failure detection.

[R-91] The pBG MUST support monitoring the status of the LSL using ARP/ND ping.
[R-92] The pBG SHOULD support monitoring the status of the LSL using BFD.

When ARP/ND keep alive is used for monitoring of the LSL:

[R-93] The pBG MUST support the IP session keep alive mechanism described in TR-146 section 6.2.4, using ARP ping for proactive monitoring of the reachability of the default gateway in the vBG.
[R-94] The pBG MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861 [33]).
[R-95] The pBG MUST support configuration of the ARP/ND keep-alive time interval, response timeout and the detection multiplier (missed heartbeat count).

When BFD is used for monitoring of the LSL:

[R-96] The pBG MUST support single-hop BFD per RFC 5880 [37] mandatory features.
[R-97] The pBG MUST support asynchronous mode.
[R-98] The pBG MUST support a single BFD session per LSL, to the vBG IP address over the LSL.

6.3.1.1.2 Overlay LSL tunnel endpoint monitoring requirements

For Overlay LSL, the pBG should be capable of monitoring the status of the tunnel endpoint in the vBG_MUX, using ICMP. Optionally, to minimize resource consumption, ICMP monitoring could be enabled only when the vBG is not reachable, as detected by ARP/ND or BFD.

[R-99] The pBG MUST support monitoring the status of the vBG_MUX tunnel endpoint using ICMP ping.
[R-100] The pBG MUST support configuration of the ICMP queries interval, response timeout and the detection multiplier (missed heartbeat count).
[R-101] The pBG SHOULD support enabling ICMP queries to the vBG_MUX tunnel endpoint when the vBG is not reachable.

 © The Broadband Forum. All rights reserved

6.3.1.2 Connectivity monitoring by the vBG and vBG_MUX

6.3.1.2.1 LSL monitoring requirements

The vBG monitors the state of its connectivity to the pBG via the LSL monitoring mechanism, based on ARP and/or ND. BFD can be used instead of ARP/ND for fast failure detection.

[R-102] The vBG MUST support monitoring the status of the LSL using ARP/ND ping.
[R-103] The vBG SHOULD support monitoring the status of the LSL using BFD.

When ARP/ND keep alive is used for monitoring of the LSL:

[R-104] The vBG MUST support the IP session keep alive mechanism described in TR-146 section 6.2.4, using ARP ping for proactive monitoring of the reachability of the default gateway in the pBG.
[R-105] The vBG MUST support IPv6 Network Unreachability Detection for IP Sessions (RFC 4861).
[R-106] The vBG MUST support configuration of the keep-alive time interval, response timeout and the detection multiplier (missed heartbeat count).
[R-107] The vBG SHOULD support activating the ARP ping and IPv6 NUD keep-alive mechanisms only when no user traffic has been received within a configurable period.

When BFD is used for monitoring of the LSL:

[R-108] The vBG MUST support single-hop BFD per RFC 5880 mandatory features.
[R-109] The vBG MUST support asynchronous mode.
[R-110] The vBG MUST support a single BFD session per LSL, to the pBG IP address over the LSL.

6.3.1.2.2 Overlay LSL tunnel endpoint monitoring requirements

For Overlay LSL, the vBG_MUX should be capable of responding monitoring mechanism used by the pBG, based on ICMP ping.

[R-111] The vBG_MUX MUST support responding to ICMP ping requests.

6.3.1.2.3 vBG resource management upon LSL failure

In case a failure is detected, the vBG may use this information for resource management.

[R-112] If a pBG is detected as unreachable, the vBG SHOULD clear the resources used by the associated customer branch.

 © The Broadband Forum. All rights reserved

6.3.2 LSL Failure and Protection

The pBG monitors the availability of the LSL, based on mechanisms defined in section 6.3.1. Upon detection of LSL failure, the pBG may take different actions, depending on its forwarding mode and the number of WAN interfaces it has.

6.3.2.1 Handling of LSL failure

The use of a Routed pBG forwarding mode will allow keeping intra-office LAN communications going in the case of LSL failure, if the pBG is running a local DHCP server for LAN IP addressing allocation and a local DNS resolver for LAN machine names. If the pBG is doing DHCP relay and/or the DNS server for LAN devices is always set to a network server, intra-office communications will be impacted.

In case the pBG is operating in a Bridged pBG forwarding model, the pBG may still host essential vBG functionality locally, e.g. DHCP server and DNS resolver, either as part of the main software running in the pBG or in built-in NFVI in the pBG. Alternatively, it could activate backup DHCP and DNS services upon detecting a failure of the LSL, as described in section 7.1.4/TR-317. Both solutions allow intra-office LAN connectivity to remain uninterrupted.

If the pBG functions as a Routed pBG or Routing and Bridging pBG, then:
[R-113] The pBG MUST support a local DHCP server.
[R-114] The pBG MUST support a local DNS resolver.

For a bridged pBG with NFVI, then:
[R-115] The pBG MUST support hosting a local DHCP server.
[R-116] The pBG MUST support hosting a local DNS resolver.

For a bridged pBG without NFVI, then:
[R-117] The pBG MUST support a B-DHCP server, as per section 7.1.4/TR-317.
[R-118] The pBG MUST support a B-DNS resolver, as per section 7.1.4/TR-317.

6.3.2.2 LSL protection

If keeping connection to the Internet, other branches and/or services hosted in the vBG and/or cloud is essential for the business needs of an enterprise customer, the use one of the multi-homing alternatives described in section 6.2 will be required. In such case, upon detection of an LSL failure, the pBG can re-establish the LSL over a backup access link.

If a pBG supports LSL protection, then the following requirements apply:

[R-119] The pBG MUST support re-establishing the LSL over a backup access link.

When the failed access link comes back online, and is once again capable of providing LSL connectivity, it is a matter of policy whether to re-establish the LSL over the primary access link or to keep using the backup LSL.

 © The Broadband Forum. All rights reserved

[R-120] The pBG MUST support LSL revertive mode.
[R-121] The pBG MUST support LSL non-revertive mode.
[R-122] The pBG MUST support setting of the LSL revert mode based on configuration and/or policy.

6.4 Performance Monitoring Requirements

Business services normally have SLAs. Hence support for Performance Monitoring is mandatory.

[R-123] The VBG System MUST support measuring the performance between any two sites of a given business customer (pBG to pBG), as defined in section 6.4.1.
[R-124] The VBG System MUST support measuring the performance between a site and the vBG (vBG to pBG), as defined in section 6.4.2.

TWAMP Light (TWL), an IP OAM tool described in RFC 5357, Appendix I [36], shall be used for measuring the performance between any two customer sites, as well as between a customer site and the IP edge.

[R-125] The pBG MUST support acting as a TWL Session-Reflector as per Section 6.2/TR-390 [10].
[R-126] The pBG MUST support static provisioning of the TWL Session-Reflector.

6.4.1 Monitoring pBG to pBG

TWL will be used to monitor the performance of the end-to-end service between a pair of pBG's, across the respective LSLs, the network as well as the vBG hosting infrastructure. TWL will also be used to provide performance monitoring per class of service, an important attribute to monitor business SLA.

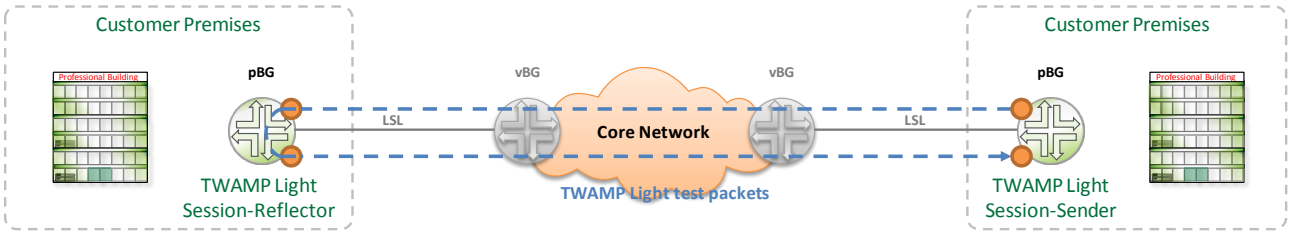

Figure 34 – pBG to pBG Performance Monitoring

[R-127] The pBG MUST support acting as a TWL Session-Sender.
[R-128] The pBG MUST be able to mark the DSCP field in the IP packet per the class of service measured.

© The Broadband Forum. All rights reserved

6.4.2 Monitoring between vBG and pBG

TWL will be used to monitor the performance of the end-to-edge service from pBG user port to the vBG, across the LSL. TWL will also provide Performance monitoring per class of service, an important attribute to monitor business SLA.

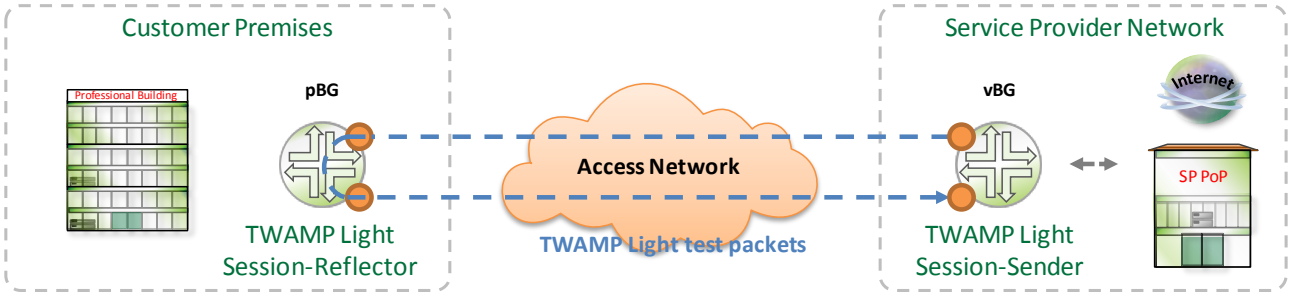

Figure 35 – vBG to pBG Performance Monitoring

[R-129] The vBG MUST support acting as a TWL Session-Sender as per Section 6.3/TR-390.
[R-130] The vBG MUST be able to mark the DSCP field in the IP packet per the class of service measured.

6.5 QoS

Business services often require a flexible Class of Service handling to allow the end customer to configure a set of criteria to meet a given CoS profile for both upstream and downstream traffic.

[R-131] The pBG MUST support QoS mechanisms on the pBG-LAN Interface, as described in Table 8
[R-132] The pBG MUST support QoS mechanisms on the pBG-LSL Interface, as described in Table 9 and Table 10
[R-133] The vBG MUST support QoS mechanisms on the vBG-LSL Interface, as described in section 6.5.2
[R-134] The vBG MUST support QoS mechanisms on the vBG-WAN Interface, as described in section 6.5.2
[R-135] QoS rules, that include classification, queuing and scheduling, MUST be configurable

6.5.1 QoS requirements on the pBG

The following requirements are based on TR-124 with additional requirements for business GW device. The following QoS requirements have been identified for the pBG-LAN interface:

Section	Item	Requirements
QOS	LAN Quality of Service requirements	
LAN.QOS	1	The pBG MUST support classification of LAN directed WAN traffic and placement into appropriate queues (or discard) based on any one or more of the following pieces of information:

© The Broadband Forum. All rights reserved

		(1) destination IP address(es) with subnet mask,
		(2) originating IP address(es) with subnet mask,
		(3) Diffserv codepoint (IETF RFC 3260 [28]),
		(4) protocol (TCP, UDP, ICMP, IGMP …),
		(5) source TCP/UDP port and port range,
		(6) destination TCP/UDP port and port range,
		(7) source MAC address,
		(8) destination MAC address,
		(9) IEEE 802.1Q Ethernet priority,
		(10) Ethertype (IEEE 802.3) length/type field), and
		(11) IEEE 802.1Q VLAN identification
LAN.QOS.	2	The pBG MUST support classification of LAN directed WAN traffic and placement into appropriate queues (or discard) based on any one or more of the following pieces of information:
		(1) source MAC address, and
		(2) destination MAC address
LAN.QOS	3	The pBG MUST support classification of LAN directed WAN traffic and placement into appropriate queues (or discard) based on any one or more of the following pieces of information:
		(1) destination IP address(es) with subnet mask,
		(2) originating IP address(es) with subnet mask,
		(3) protocol (TCP, UDP, ICMP, IGMP …),
		(4) source TCP/UDP port and port range,
		(5) destination TCP/UDP port and port range, and
		(6) IEEE 802.1Q VLAN identification
LAN.QOS	4	The pBG MUST be able to mark or remark the Diffserv codepoint of traffic identified based on any of the classifiers supported by the PBG
LAN.QOS	5	The pBG MUST support a minimum of four downstream queues per LAN port.
LAN.QOS	6	The pBG MUST duplicate the set of queues for each LAN egress port. This can be done logically or physically.
LAN.QOS	7	The pBG SHOULD be able to configure each queue for strict priority or weighted round robin scheduling. Strict priority queues are served with priority over all other queues. WRR queues are served on the basis of configurable weights.
LAN.QOS	8	The WRR Queues should support congestion avoidance mechanism, e.g., WRED, RED.
LAN.QOS	9	The pBG MUST provide counters in terms of dropped and emitted packets/bytes for each queue. Statistics SHOULD be from the time last counter or on a configurable sample interval
LAN.QOS	10	The pBG SHOULD provide counters in terms of dropped and emitted packets/bytes for each VPN/VRF (VLAN) and queue. Statistics SHOULD be from the time last counter or on a configurable sample interval
	11	The pBG MUST provide information about queue occupancy in terms

 © The Broadband Forum. All rights reserved

LAN.QOS		of packets and peak percentage. Statistics SHOULD be collected from the time of last counter reset or on a configurable sample interval.
LAN.QOS	12	The pBG SHOULD be able to monitor the physical layer rate of the LAN interfaces, maintaining information about the current available bandwidth and measurement history.

Table 8 – pBG-LAN interface QoS requirements

The following QoS requirements have been identified for the pBG-LSL interface:

Section	Item	Requirements
QOS	WAN Quality of Service requirements	
WAN.QoS.	1	The pBG MUST support classification of WAN directed LAN traffic or LAN directed LAN traffic (local switching), and placement into appropriate queues (or discard) based on any one or more of the following pieces of information:
		(1) destination IP (v4 or v6) address(es) with subnet mask,
		(2) originating IP (v4 or v6) address(es) with subnet mask,
		(3) source MAC address,
		(4) destination MAC address,
		(5) protocol (TCP, UDP, ICMP, IGMP, …),
		(6) source TCP/UDP port and port range,
		(7) destination TCP/UDP port and port range,
		(8) IEEE 802.1Q Ethernet priority,
		(9) FQDN (fully qualified domain name) of WAN session,
		(10) Diffserv codepoint (IETF RFC 3260),
		(11) Ethertype (IEEE 802.3) length/type field),
		(12) traffic handled by an ALG,
		(13) IEEE 802.1Q VLAN identification,
		(14) Wi-Fi SSID and,
		(15) LAN type (Ethernet, Wi-Fi, etc.)
WAN.QOS.	2	The pBG MUST support classification of WAN directed LAN traffic or LAN directed LAN traffic (local switching), and placement into appropriate queues (or discard) based on any one or more of the following pieces of information:
		(1) destination IP address(es) with subnet mask,
		(2) originating IP address(es) with subnet mask,
		(3) protocol (TCP, UDP, ICMP, …),
		(4) source TCP/UDP port and port range,
		(5) destination TCP/UDP port and port range, and
		(6) VLAN ID
WAN.QOS	3	The pBG MUST support the differentiated services field (DS field) in IP (v4 or v6) headers as defined in IETF RFC 2474 [21].
WAN.QOS	4	The pBG MUST by default recognize and provide appropriate treatment to packets marked with recommended Diffserv code points whose values and

 © The Broadband Forum. All rights reserved

<table>
<tr><td></td><td></td><td colspan="4">behavior are defined in IETF RFCs 2474, 2475 [22], 2597 [23], 3246 [27], and 3260. Specifically, the values shown in the DSCP column of the table below MUST be supported, except Cs0-7, which are optional.</td></tr>
<tr><td></td><td></td><td>Class</td><td>Description</td><td>DSCP marking (name)</td><td>DSCP marking (decimal value)</td></tr>
<tr><td></td><td></td><td>EF</td><td>Realtime</td><td>ef</td><td>46</td></tr>
<tr><td></td><td></td><td>AF4 – in-contract</td><td>Premium class4 (in)</td><td>af41</td><td>34</td></tr>
<tr><td></td><td></td><td>AF4 – out-of-contract</td><td>Premium class4 (out)</td><td>af42, af43</td><td>36, 38</td></tr>
<tr><td></td><td></td><td>AF3 – in-contract</td><td>Premium class3 (in)</td><td>af31</td><td>26</td></tr>
<tr><td></td><td></td><td>AF3 – out-of-contract</td><td>Premium class3 (out)</td><td>af32, af33</td><td>28, 30</td></tr>
<tr><td></td><td></td><td>AF2 – in-contract</td><td>Premium class2 (in)</td><td>af21</td><td>18</td></tr>
<tr><td></td><td></td><td>AF2 – out-of-contract</td><td>Premium class2 (out)</td><td>af22, af23</td><td>20, 22</td></tr>
<tr><td></td><td></td><td>AF1 – in-contract</td><td>Premium class1 (in)</td><td>af11</td><td>10</td></tr>
<tr><td></td><td></td><td>AF1 – out-of-contract</td><td>Premium class1 (out)</td><td>af12, af13</td><td>12, 14</td></tr>
<tr><td></td><td></td><td>DE/BE</td><td>Default / Best Effort</td><td>be</td><td>0</td></tr>
<tr><td></td><td></td><td>Cs0 (optional)</td><td>Class Selector 0</td><td>cs0</td><td>0</td></tr>
<tr><td></td><td></td><td>Cs1 (optional)</td><td>Class Selector 1</td><td>cs1</td><td>8</td></tr>
<tr><td></td><td></td><td>Cs2 (optional)</td><td>Class Selector 2</td><td>cs2</td><td>16</td></tr>
<tr><td></td><td></td><td>Cs3 (optional)</td><td>Class Selector 3</td><td>cs3</td><td>24</td></tr>
<tr><td></td><td></td><td>Cs4 (optional)</td><td>Class Selector 4</td><td>cs4</td><td>32</td></tr>
<tr><td></td><td></td><td>Cs5 (optional)</td><td>Class Selector 5</td><td>cs5</td><td>40</td></tr>
<tr><td></td><td></td><td>Cs6 (optional)</td><td>Class Selector 6</td><td>cs6</td><td>48</td></tr>
<tr><td></td><td></td><td>Cs7 (optional)</td><td>Class Selector 7</td><td>cs7</td><td>56</td></tr>
<tr><td>WAN.QOS.</td><td>6</td><td colspan="4">The pBG MUST be able to mark or remark the Diffserv codepoint or IEEE 802.1Q Ethernet priority of traffic identified based on any of the classifiers supported by the pBG.</td></tr>
<tr><td>WAN.QOS.</td><td>7</td><td colspan="4">The pBG MUST support one best effort (BE) queue, one expedited forwarding (EF) queue and a minimum of four assured forwarding (AF) queues.</td></tr>
<tr><td>WAN.QOS.</td><td>8</td><td colspan="4">The pBG MUST duplicate the set of queues for each access session (e.g., VLAN). This can be done logically or physically.</td></tr>
<tr><td>WAN.QOS.</td><td>9</td><td colspan="4">The pBG MUST support the appropriate mechanism to effectively implement Diffserv per-hop scheduling behaviors.

The pBG MUST be able to configure each queue defined in WAN.QoS.7 for strict priority or weighted round robin scheduling. Strict priority queues are served with priority over all other queues. A strict priority scheduler is preferred for EF. WRR queues are served on the basis of configurable weights, provided with a mechanism to prevent starvation (WRR queue minimum bandwidth).</td></tr>
<tr><td>WAN.QOS.</td><td>10</td><td colspan="4">The WRR Queues should support congestion avoidance mechanism, e.g., WRED, RED.</td></tr>
<tr><td>WAN.QOS.</td><td>11</td><td colspan="4">The pBG MUST support aggregate shaping of upstream traffic across all access sessions (e.g. VLAN)</td></tr>
</table>

 © The Broadband Forum. All rights reserved

WAN.QOS.	12	The pBG SHOULD support aggregate shaping of upstream traffic per VRF/VPN (VLAN)
WAN.QOS.	13	The pBG MUST support per-class shaping of upstream traffic. Classes are defined in WAN.QoS.4.
WAN.QOS.	14	The pBG MUST support the capability to fragment IP traffic on sessions that it originates, in order to limit the effect of large packets on traffic delay.
WAN.QOS.	15	The packet size threshold before fragmenting AF and BE packets MUST be configurable.
WAN.QOS.	16	The pBG MUST handle all telephone service-related network traffic by a high priority queue to avoid congestion, delay, jitter, or packet loss.
WAN.QOS.	17	The pBG MAY handle all telephone service-related network traffic by a dedicated WAN interface to avoid congestion, delay, jitter, or packet loss.
WAN.QOS.	18	The pBG MUST provide counters in terms of dropped and emitted packets/bytes for each queue. Statistics SHOULD be collected from the time of last counter reset or on a configurable sample interval.
WAN.QOS.	19	The pBG SHOULD provide counters in terms of dropped and emitted packets/bytes for each VPN/VRF(VLAN)/queue. Statistics SHOULD be collected from the time of last counter reset or on a configurable sample interval.
WAN.QOS.	20	The pBG MUST provide information about queue occupancy in terms of packets and peak percentage. Statistics SHOULD be collected from the time of last counter reset or on a configurable sample interval.
WAN.QOS.	21	The pBG MUST support classification of WAN-directed internally-generated traffic and placement into appropriate queues based on any one or more of the following pieces of information:
		(1) destination IP address(es) with subnet mask,
		(2) originating IP address(es) with subnet mask,
		(3) protocol (TCP, UDP, ICMP, …),
		(4) source TCP/UDP port and port range,
		(5) destination TCP/UDP port and port range,
		(6) Diffserv codepoint (IETF RFC 3260), and
		(7) physical port, in case of voice packets.
WAN.QOS.	22	Classification Key configured by external protocol – TBD

Table 9 – pBG-LSL interface QoS requirements

In addition, the following requirements have been identified for the pBG-LSL interface, when using VLANs:

Section	Item	Requirements
QOS	VLAN Quality of Service requirements	
WAN.QoS.VLAN.	1	The pBG MUST support sending the following frame types: Untagged frames, priority-tagged frames, and VLAN-tagged frames and double-tagged in the upstream direction.

© The Broadband Forum. All rights reserved

WAN.QoS.VLAN	2	The pBG MUST support setting the priority tag and VLAN ID Values both at the Inner tag and Outer tag per classification, or per DSCP value.
WAN.QoS.VLAN	3	The pBG MUST support receiving untagged, VLAN-tagged and double tagged Ethernet frames in the downstream direction, and MUST be able to strip the VLAN Tagging from the one received tagged and double tagged (per configuration).

Table 10 – pBG-LSL interface VLAN-related QoS requirements

6.5.2 QoS requirements on the vBG

In current networks, a BG for SMB/SOHO type of customers typically has a single service and can be served by an MS-BNG. More complex customer branches will normally have a multi-service BG, e.g. Internet Access, VPN, L2 service, etc., served by MS-BNG or a business Provider Edge (PE) router. The VBG System architecture is aimed at enhancing all these scenarios. Note however, that even if there are architectural changes when deploying the VBG System, that the QoS requirements largely remain the same as for existing services, with the addition of the mechanisms required for the Overlay LSL case.

In the scenario where the vBG is hosted at the MS-BNG's QoS is implemented within the MS-BNG, the MS-BNG hierarchical scheduling function can perform hierarchical scheduling as per the requirements documented in section 4.4.3/TR-178 (Hierarchical QoS) and section 7.1.4/TR-178 (Hierarchical QoS on MS-BNG).

In the scenario where the vBG is hosted outside of the MS-BNG, the vBG needs to include a Traffic Management Function (TMF) that provides at a minimum an aggregate shaping/policing rate limit for upstream and downstream traffic per pBG, as well as differential queuing.

6.5.2.1 Classification requirements

[R-136] The vBG MUST support ingress traffic classification using 802.1p, IP precedence and DSCP on the vBG-LSL and vBG-WAN interfaces

[R-137] The vBG MUST support ingress traffic classification using multi-field criteria, supporting at least the following header fields:

- Source/Destination IP address
- Source/Destination port
- Protocol number
- DSCP
- 802.1p bits
- VLAN ID

[R-138] The vBG MUST support ingress remarking of the DSCP bits, based on the resulting forwarding class.

[R-139] The vBG MUST support egress marking of the 802.1p bits, based on the forwarding class of the traffic.

[R-140] The vBG MUST support egress remarking of the DSCP bits, based on the forwarding class of the traffic.

 © The Broadband Forum. All rights reserved

In addition, following requirements apply to the vBG_MUX in the Overlay LSL mode:

[R-141] The vBG_MUX MUST support ingress traffic classification using 802.1p, IP precedence and DSCP of the customer traffic coming from the vBG.
[R-142] The vBG_MUX MUST support ingress traffic classification using 802.1p, IP precedence and DSCP of the LSL tunnel header of traffic coming from the pBG.
[R-143] The vBG_MUX MUST support egress marking of 802.1p and DSCP priority bits of the LSL tunnel header on the interface to the pBG.

6.5.2.2 Queueing and policing requirements

[R-144] The vBG MUST support egress queuing per pBG and per traffic class on the vBG-LSL interface.
[R-145] The vBG SHOULD support ingress queuing per pBG and per traffic class on the vBG-LSL interface.
[R-146] The vBG SHOULD support egress policing per pBG and per traffic class on the vBG-LSL interface.
[R-147] The vBG MUST support ingress policing per pBG and per traffic class on the vBG-LSL interface.
[R-148] The vBG MUST support ingress queuing per traffic class on the vBG-WAN interface.
[R-149] The vBG MUST support egress queuing per traffic class on the vBG-WAN interface.
[R-150] The vBG MUST be able to assign a minimum bandwidth guarantee to a queue or policer. **Note:** the minimum can be zero (no guarantee).

6.5.2.3 Scheduling requirements

[R-151] The vBG hierarchical scheduler MUST support the following types of scheduling: strict priority, weighted round robin, and round robin.
[R-152] The vBG MUST be able to provide an egress aggregate shaping/policing rate per LSL for all queues and policers.
[R-153] The vBG SHOULD be able to provide an ingress aggregate shaping/policing rate per LSL for all queues and policers.
[R-154] The vBG MUST support configuring a queue with a given forwarding behavior such as Expedited Forwarding, Assured Forwarding, etc.

 © The Broadband Forum. All rights reserved

6.6 IP addressing

6.6.1 Address assignment

In the VBG System architecture, IP address assignment to end-devices in the customer branches may be performed by different functions, depending on the deployment model:

pBG mode	IP address family	pBG role	vBG role
Bridged	IPv4	None	DHCP Server
		DHCP Server	None
	IPv6	None	DHCPv6 Server
		DHCPv6 Server	SLAAC RAs
Routed	IPv4	DHCP Relay	DHCP Server
		DHCP Server	None
	IPv6	DHCPv6 Relay SLAAC RAs	DHCPv6 Server
		DHCPv6 Server SLAAC RAs	None

Table 11 – VBG System IP address assignment roles

This section provides the requirements on IP assignment for both IPv4 and IPv6.The pBG and vBG must comply with the requirements listed in the following sub-sections, where the word "Device" must be read as "pBG" or "vBG", depending on the roles performed by the pBG and vBG (see Table 11).

6.6.1.1 IPv4 addressing

Customer branches within the same enterprise customer will have non-overlapping subnets (one or multiple subnets per site). However, IP subnets of different enterprise customers will typically overlap. Each LAN device will obtain a unique private address from the one of the IP subnets in the branch they are connected to. The private addresses are then translated via the vBG NAT function to provide Internet access.

For Devices performing a DHCP Server function, the following requirements need to be met:

[R-155] The Device MUST support a DHCP Server that is compliant with the TR-124 LAN.DHCPS (2 4, 7, 11, 12, 13, 16, 17, 18) requirements.

[R-156] The Device MUST support a DHCP Server that is compliant with the standards below and required elements of associated updates:

- RFC 2131 Dynamic Host Configuration Protocol [18]
- RFC 2132 DHCP Options and BOOTP Vendor Extensions [19]
- RFC 2939 Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types [25]

 © The Broadband Forum. All rights reserved

For Devices performing a DHCP Relay function, the following requirements need to be met:

[R-157] The Device MUST support a DHCP Relay Agent function as described in RFC 951 “BOOTP” [14], RFC 2131 “DHCP” and RFC 3046 “DHCP Relay Agent Information Option” [26] on the LAN interfaces.

6.6.1.2 IPv6 addressing

IPv6 address assignment to end-devices can be done by means of DHCPv6 or Stateless Address Auto-configuration (SLAAC). To support DHCPv6, the Device must satisfy the following requirements:

For Devices performing a DHCPv6 Server function, the following requirements need to be met:

[R-158] The Device MUST support a DHCPv6 Server that is compliant with the TR-124 LAN.DHCPv6S (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12) requirements.

For Devices performing a DHCPv6 Relay function, the following requirements need to be met:

[R-159] The Device MUST support a DHCPv6 Relay Agent function as described in RFC 3315 [29].

For Devices performing a SLAAC Router function, the requirements below need to be met. Note that unlike DHCP leases, SLAAC address allocations are not stateful and it is difficult to track the connectivity of the end-device.

[R-160] The Device MUST be able to advertise (RA) a prefix per LAN segment.
[R-161] The Device MUST support Neighbor Unreachability Detection (NUD), as defined in RFC 4861, “Neighbor Discovery in IPv6”, to verify if an end-device is still connected to the network.

6.6.2 NA(P)T

[R-162] The vBG MUST support Source NAT
[R-163] The vBG MUST support Source NAT with Port Address Translation (PAT)
[R-164] The vBG MUST support Destination NAT with PAT
[R-165] The vBG MUST support NAT static port forwarding
[R-166] The vBG MUST support NAT64

6.7 Forwarding and routing protocols

6.7.1 pBG Ethernet forwarding requirements

If the pBG functions as a Bridged pBG or Routing and Bridging pBG, as per section 0, then:

[R-167] The pBG bridge MUST conform to 802.1Q-2014 and consequently support MAC learning.

 © The Broadband Forum. All rights reserved

6.7.2 pBG IP forwarding requirements

If the pBG functions as a Routed pBG or Routing and Bridging pBG, as per section 0, then:

[R-168] The pBG MUST support IPv4/IPv6 dual stack functionality.
[R-169] The pBG MUST support forwarding IPv4 traffic.
[R-170] The pBG MUST support forwarding IPv6 traffic.
[R-171] The pBG MUST support requirement Gen.Net.4/TR-124

6.7.3 vBG IP forwarding requirements

[R-172] The vBG MUST support IPv4/IPv6 dual stack functionality.
[R-173] The vBG MUST support forwarding IPv4traffic.
[R-174] The vBG MUST support forwarding IPv6 traffic.

6.7.4 pBG routing and protocols requirements

The following requirements apply in the case the pBG functions as Routed pBG or Routing and Bridging pBG, as per section 0.

The following protocols apply to the pBG-LSL interface:

[R-175] The pBG MUST support Static routes
[R-176] The pBG MUST support BGP-4, as per RFC 4271 [31].

The following protocols apply to the pBG-LAN interface:

[R-177] The pBG MUST support Static routes
[R-178] The pBG SHOULD support BGP-4, as per RFC 4271.
[R-179] The pBG SHOULD support OSPFv2, as per RFC 2328 [20].
[R-180] The pBG SHOULD support OSPFv3, as per RFC 5340 [35].
[R-181] The pBG SHOULD support IS-IS, as per RFC 1195 [15].

6.7.5 vBG routing and protocols requirements

The following requirements apply to the LSL interface:

[R-182] The pBG MUST support Static routes
[R-183] The pBG MUST support BGP-4, as per RFC 4271.

The following requirements apply to the WAN interface:

[R-184] The vBG MUST support Static routes
[R-185] The vBG MUST support OSPFv2, as per RFC 2328.
[R-186] The vBG MUST support OSPFv3, as per RFC 5340.
[R-187] The vBG MUST support BGP-4, as per RFC 4271.
[R-188] The vBG MUST support IS-IS, as per RFC 1195.

 © The Broadband Forum. All rights reserved

6.8 Security

If the vBG supports security:

[R-189] The vBG MUST support Firewall
[R-190] The vBG MUST support IPsec VPN on the WAN side
[R-191] The vBG MUST support Tunnels (GRE, IP-IP) on the WAN side
[R-192] The vBG SHOULD support Network attack detection
[R-193] The vBG SHOULD support DDoS prevention

For IPsec VPNs on the WAN side, the following requirements apply:

[R-194] The vBG MUST support encryption using Advanced Encryption Standard (AES-CBC), using a minimum key size of 128 bits
[R-195] The vBG SHOULD support encryption using triple Data Encryption Standard (3DES-CBC)
[R-196] The vBG SHOULD NOT support encryption using Data Encryption Standard (DES-CBC)
[R-197] The vBG MUST support authentication using SHA-256
[R-198] The vBG MUST support authentication using SHA-128
[R-199] The vBG SHOULD support authentication using SHA-1
[R-200] The vBG SHOULD NOT support authentication using Message Digest 5 (MD5)

6.9 AAA requirements

6.9.1 Flat LSL setup

Similar to the mechanisms described in TR-317, in the case where the vBG functions are not located at the MS-BNG, the MS-BNG needs to create connectivity from the pBG to the vBG_MUX. Consequently, the MS-BNG needs to be able to:

- Extend the subscriber's access VLAN to the vBG_MUX, in case of Flat Ethernet LSL
- Provide the pBG with tunneling information, in the case of Overlay Ethernet LSL

Section 7.1.3.1/TR-317 describes RADIUS attributes that can be sent by the AAA to the MS-BNG to assist setting up the connectivity between the pBG and the vBG_MUX. Section 7.1.3.2/TR-317 describes the use of these RADIUS attributes and associated requirements for the Flat LSL case, while Section 7.1.3.3/TR-317 describes its use and requirements for the Overlay LSL case.

[R-201] The MS-BNG MUST conform to Section 7.1.3.2/TR-317.
[R-202] The MS-BNG MUST conform to Section 7.1.3.3/TR-317.

In turn, the vBG_MUX needs to know to which vBG the user traffic coming from a given pBG must be forwarded to. Section 7.1.3.4/TR-317 describes mechanisms and requirements that are applicable for the dynamic setup of pBG to vBG associations at the vBG_MUX.

[R-203] The vBG_MUX MUST support the requirements in Section 7.1.3.4.1/TR-317.
[R-204] The vBG_MUX MUST support the requirements in Section 7.1.3.4.2/TR-317.

 © The Broadband Forum. All rights reserved

6.9.2 Overlay LSL authentication

When the Overlay LSL connection between the pBG and the vBG is set up over the network of the VBG System service provider, implicit authentication of the LSL can be assumed. The pBG underlay connection will be authenticated with the usual means in the MSBN:

- 802.1X [12] access line authentication, with the AN acting as Authenticator
- IP session authentication, with the MS-BNG validating the access line credentials with a AAA server (e.g. DHCP Option 82 Circuit-ID, PPPoE PAP/CHAP credentials, etc.)

In addition, after the IP session has been authenticated, the MS-BNG performs anti-spoofing of the MAC/IP pair associated with the IP session.

However, when one or more of the pBG's access links are connected over a network service provider different than the VBG System service provider, it may be desirable to authenticate the pBG prior to allowing traffic over the Overlay LSL.

For basic authentication, the pBG authenticates itself to a AAA server of the VBG System service provider using EAP. That is, the pBG performs the roles of Supplicant and Authenticator in 802.1X, while the AAA server acts as the Authentication Server. The vBG_MUX could act as a RADIUS proxy, caching the authentication state of the pBG and its source IP address.

6.9.2.1 pBG requirements

[R-205] The pBG MUST be able to authenticate itself using a RADIUS client to a fixed access AAA server, collapsing the functions of IEEE 802.1X supplicant and authenticator.

[R-206] The pBG MUST support Extensible Authentication Protocol-Transport Layer Security (EAP-TLS)

[R-207] The pBG MUST support Protected EAP (EAP-PEAP) with MS-CHAPv2 username and password credentials

[R-208] Each pBG MUST have a unique factory-installed private/public key pair and an embedded ITU-T X.509 version 3 / IETF RFC 5280 [34] certificate that has been signed by the pBG vendor's certificate authority

[R-209] The pBG certificate MUST have a validity period greater than the operational lifetime of the pBG

6.9.2.2 vBG_MUX requirements

[R-210] The vBG_MUX MUST support a RADIUS proxy to relay EAP packets to the AAA server.

[R-211] The vBG_MUX MUST cache the IP of the pBG as a valid LSL source, learned as part of the AAA process.

6.9.3 Using RADIUS AAA to dynamically provision business services

The vBG provides access to one or more business services to an enterprise customer. To simplify the overall process, service provisioning can be automated by means of interaction between the vBG and an AAA server.

 © The Broadband Forum. All rights reserved

The AAA server will contain a list of service attributes that will be used for automated instantiation of VBG System services in the vBG. Each enterprise customer could have their slight variation of these "a la carte" attributes, using a sub-set of these to create their VBG System service(s). The AAA attributes are used by the vBG, together with service templates, to instantiate a service.

When an enterprise pBG boots up, during the setup of the pBG IP Session (e.g. DHCP, data-triggered) the vBG can trigger an AAA operation, and retrieve the service attributes of one or more enterprise services. Based on the returned attributes from AAA, the vBG automatically instantiates one or more services for the enterprise customer. Authentication can be based on the line attributes, pBG attributes (e.g. MAC address), or even the pBG IP address, in the case of data-triggered services.

There are numerous benefits of such model listed below:

- It alleviates the dependency on an external system such as OSS for service provisioning.
- It offers a dynamic service provisioning model, matching a constantly evolving access network.
- Seamless migration of business services to a different vBG
- More generically, it provides the possibility of mobile business service model in which the enterprise customer can relocate and have their services instantly created.

In summary, using AAA to dynamically provision business services eliminates lead-time and minimizes pre-planning work.

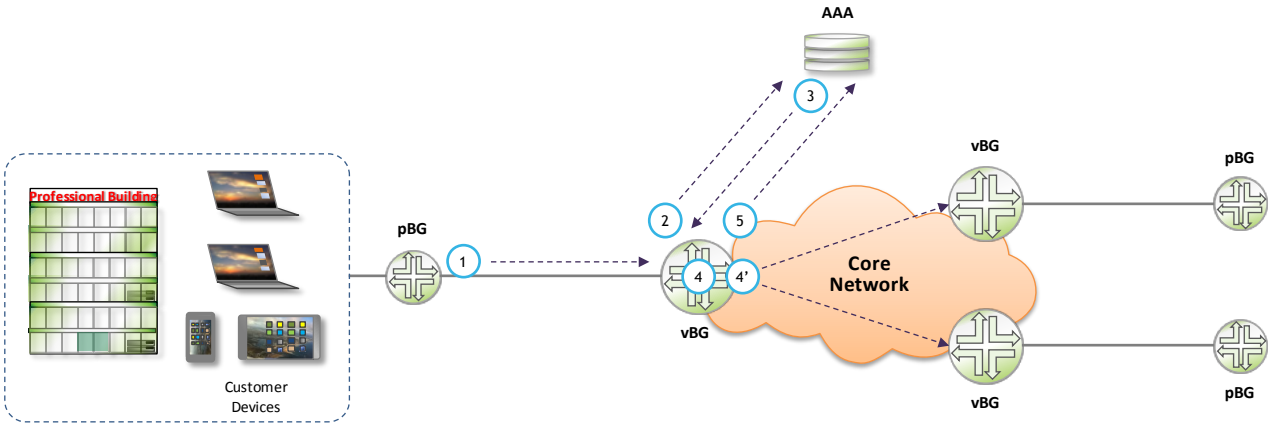

Figure 36 – AAA-based Dynamic VBG System services

Two possible service instantiation models are supported in TR-328:

- **Control Channel:** Using a dedicated VLAN to serve as a control channel for the pBG. Separate VLANs are used to transport the VBG System business services for the enterprise customer. The IP session over the control channel VLAN can be used for management and health check of the pBG, in addition to being used as a trigger to set up the VBG System business services in the vBG.

© The Broadband Forum. All rights reserved

- **Data Triggered:** The first IP packet in a VBG System service VLAN is used as a trigger for authentication and health check of the pBG. Essentially, the data channel also becomes the control channel. The AAA server can return service attributes to the vBG for that VLAN, and any other VLANs in use for VBG System services. After a successful authentication, the pBG can begin sending traffic to the vBG

As described in section 6.1, there are two different LSL transport options: Flat LSL and Overlay LSL. The operation of both RADIUS-based automated service instantiation models, Control Channel and Data Triggered, is the same for both LSL transport options and is described below.

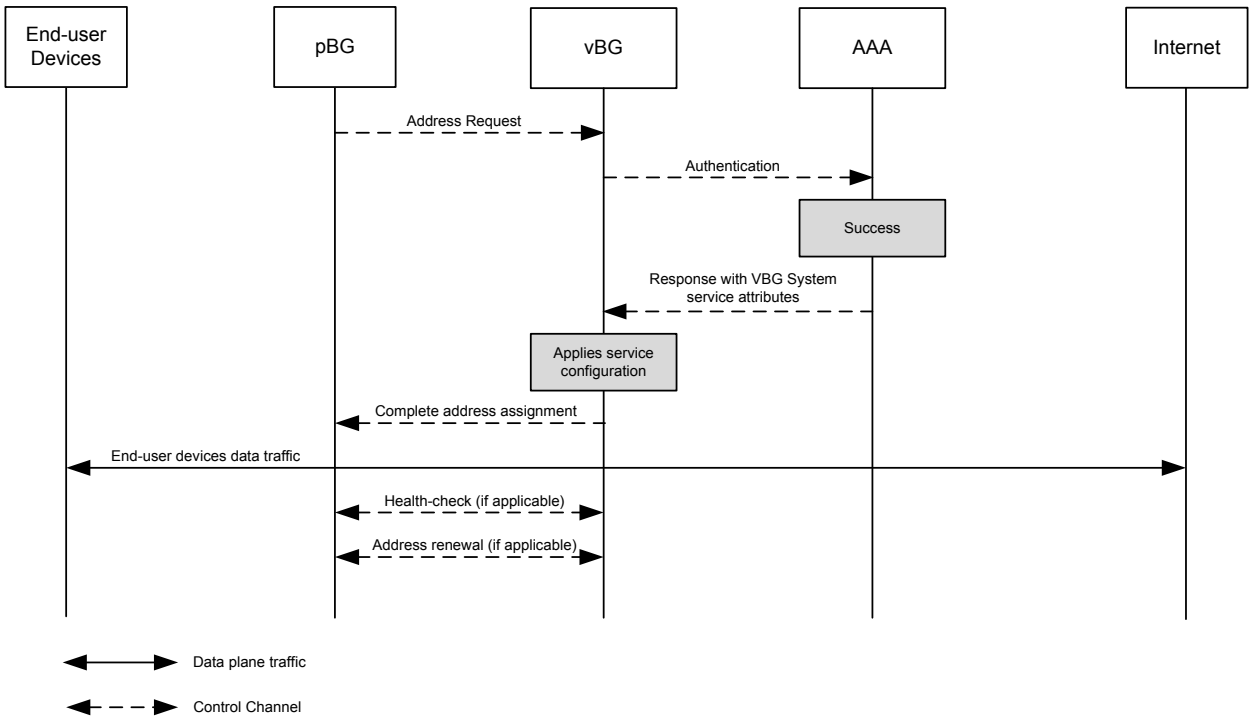

Figure 37 – Automated VBG System services – Control Channel model

Figure 37 describes the call flow for the automated instantiation of VBG System services at the vBG, using a dedicated control channel with the pBG:

1. The pBG initiates an address request (e.g. DHCP) on the control channel (VLAN) over its pBG-LSL interface. The vBG receives the request on its vBG-LSL interface
2. The vBG sends a RADIUS Access-Request to the AAA server, with the pBG's credentials. The credentials could include the Circuit ID, Interface ID, Remote ID, NAS Port ID, etc.
3. The AAA server returns a set of service attributes, which triggers the automated instantiation of one or more VBG System services in the vBG, over separate VLANs from the one used for the control channel IP session. Service attributes returned by the AAA server may include: service type, QoS parameters, filter rules, etc.
4. After completing the instantiation of the VBG System services for the enterprise customer, the vBG completes address assignment for the pBG via the control channel.
5. When the pBG receives the address assignment and the control channel IP session setup is complete, it can begin sending traffic over the VBG System service VLAN(s).

 © The Broadband Forum. All rights reserved

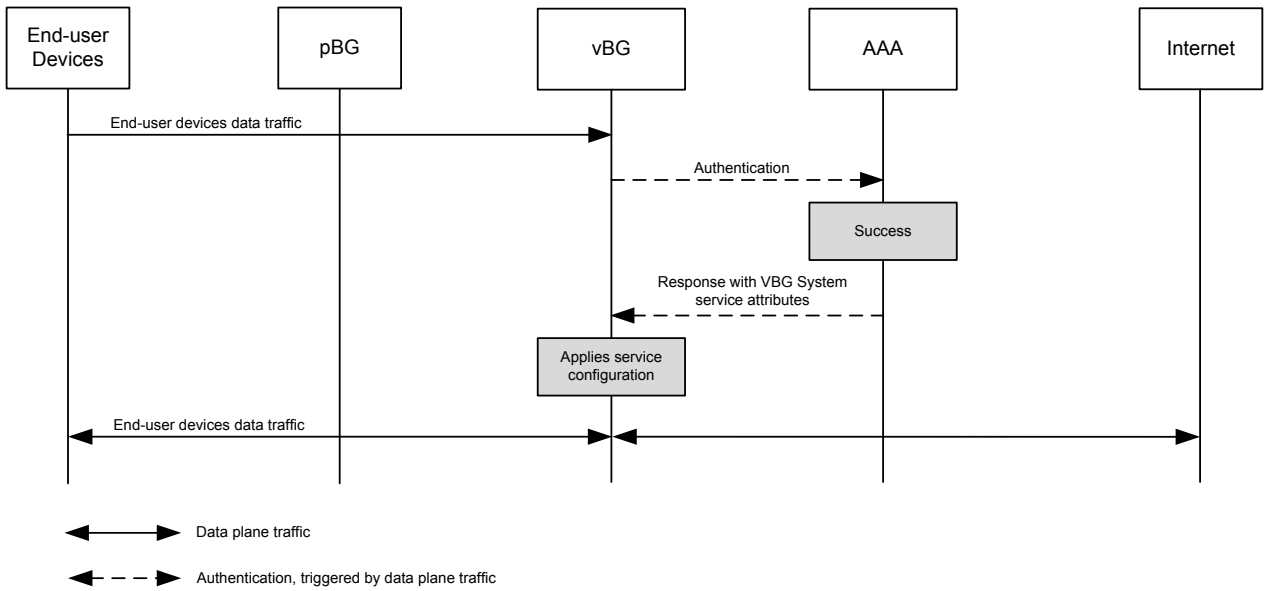

Figure 38 – Automated VBG System services – Data Triggered model

Figure 38 describes the call flow for the automated instantiation of VBG System services at the vBG, using a data triggered model over a VBG System service VLAN:

1. The pBG boots up and starts sending traffic to the vBG. The vBG receives the traffic on its vBG-LSL interface
2. Upon reception of the first packet, the vBG sends a RADIUS Access-Request to the AAA server, with the pBG's credentials. The pBG credentials could include the data in the headers of the received packet, such as source MAC address, source IP address, VLAN, NAS Port ID, etc.
3. The AAA server returns a set of service attributes, which triggers the automated instantiation of one or more VBG System services in the vBG, including at least a service for the VLAN used for the data trigger. Service attributes returned by the AAA server may include: service type, QoS parameters, filter rules, etc.
4. Once the pBG has been successfully authenticated, all traffic from that pBG is now allowed through the for the VBG System services at the vBG.

Note that the use of the Control Channel model, as well as multiple services for the Data Triggered model both require multiple VLANs on the LSL interface. Section 0 describes the operation and requirements for a multi-VLAN LSL.

In the Control Channel model, the lifetime of the VBG System services is tied to the lifetime of the control channel IP session. For the Data Triggered case, the lifetime of the services is tied to the availability of the LSL, using the monitoring mechanisms described in section 6.3. In both cases, at the end of the session, the VBG System services are deleted from the vBG and a RADIUS Accounting-Stop message is sent to the AAA server.

 © The Broadband Forum. All rights reserved

The following are the requirements for support of AAA-driven automated instantiation of VBG System services:

[R-212] The vBG MUST support control channel IP sessions
[R-213] The vBG MUST support instantiating one or more business services in different VLANs from the control channel, using service attributes received from RADIUS
[R-214] The vBG MUST support data triggered IP sessions
[R-215] The vBG MUST support instantiating one or more business services based on the setup of a data triggered IP session, using service attributes received from RADIUS.
[R-216] The vBG MUST support dynamic changes of the attributes of a business service, using RADIUS Change of Authorization (CoA) messages
[R-217] The vBG MUST support deleting a business service, using RADIUS Disconnect messages
[R-218] The vBG MUST support RADIUS accounting for the dynamically created business services
[R-219] The vBG MUST send a RADIUS Accounting-Stop message to the AAA server at the end of a control channel or data triggered session for a pBG
[R-220] The pBG MUST support setting up a control channel IP session to the vBG, over a dedicated VLAN in the pBG-LSL interface
[R-221] The pBG MUST support tying the status of one or more VBG System services to the status of a control channel IP session

The multi-VLAN LSL requirements in section 0 are also applicable here.

6.10 VBG System functional requirements

6.10.1 Backward compatibility

In order to deploy seamlessly such architecture, the backward compatibility must be preserved.

[R-222] The VBG System MUST support IP/MPLS VPN services.
[R-223] The VBG System MUST support the coexistence of both vBG(s) and legacy BG(s) on a given customer's VPN.
[R-224] The VBG System MUST support both direct Internet access and VPNs services.

6.10.2 Virtualization of Compute resources

Following are the requirements for the VBG System regarding VNF placement, mode of operation and NF service chaining:

[R-225] The VBG System MUST support VNF in 1:N mode as defined in ETSI NFV-INF 001 V0.3.8.
[R-226] The VBG System MUST support placement of VNFs on either side of the LSL, when the pBG contains NFVI node.
[R-227] The VBG System MUST support service function chaining.

 © The Broadband Forum. All rights reserved

6.11 Management

This document does not fully define the internal architecture of the VBG System and the functional distribution, so only the management of the VBG System functions is described here. Depending on the VBG System implementation, other management and orchestration aspects may be necessary, such as VNF provisioning and lifecycle management, SDN based service chaining, etc. These are out of scope for this document.

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

The network components that compose the VBG System have to be managed by the service provider.

[R-228] The vBG MUST be manageable by the service provider.
[R-229] The pBG MUST be manageable by the service provider.

6.11.1 pBG Management Client

The pBG's Management Client provides the capabilities to configure and monitor the pBG. TR-124 provides a set of management requirements for the functionality provided by the pBG, as described in the sections titled MGMT.GEN and MGMT.REMOTE.TR-069.

[R-230] The pBG MUST comply with the TR-124 MGMT.GEN requirements.
[R-231] The pBG SHOULD support being managed using SNMP and/or NETCONF
[R-232] The pBG SHOULD support being managed using TR-069
[R-233] If the pBG supports management using TR-069, the pBG Management Client MUST be compliant with the TR-124 MGMT.REMOTE.TR-069 requirements.

Management of the pBG can be done through the LSL or over a separate interface. As described in section 0, the LSL may be set to use multiple VLANs for several purposes. In such case, the pBG can be managed over one of the VLANs used for VBG System services, or over a dedicated VLAN for management purposes.

[R-234] The pBG MUST support in-band management over a VBG System service
[R-235] The pBG MUST support management over a dedicated VLAN in the pBG-LSL interface

In the case of Overlay LSL, management may be done using the pBG WAN interface (underlay) or through the Overlay pBG-LSL, either through one of the VBG System services, or a separate VLAN for management purposes. In addition to the above requirements, for Overlay LSL the following requirements apply to the pBG:

[R-236] The pBG MUST support in-band management over pBG WAN interface (underlay)

 © The Broadband Forum. All rights reserved

Note that if management of the pBG is done through the LSL and the vBG is performing NAT functions, that it will be required to open a port in the vBG NAT function for the management systems to be able to reach the pBG. This may be done statically in the vBG, as per [R-165] in section 6.6.2, or it may be triggered dynamically by the pBG, using Port Control Protocol (PCP).

6.11.2 vBG Management Client

The vBG Management Client provides fault, configuration and performance monitoring capabilities and management of application layer functions. Depending on the distribution of vBG Network Functions, a vBG can have multiple vBG Management Clients, where a vBG Management Client manages one or more of those Network Functions.

On top of that, a vBG Management Client can manage functionality for one or more vBG's. For instance, a vBG Management Client can be hosted within a vBG Hosting infrastructure (e.g., hosted by the MS-BNG), providing multi-tenant management, or it can be hosted as a function within the vBG, providing management for single business customer.

The protocol and information elements (e.g., SNMP, NETCONF/YANG) used by the vBG Management Client are beyond the scope of this document.

[R-237] For each Network Function of the vBG, the vBG MUST support a vBG Management Client to manage the fault, configuration, performance monitoring capabilities and management of the function's application layer.

End of Broadband Forum Technical Report TR-328

 © The Broadband Forum. All rights reserved