PacketCable™

Enterprise SIP Gateway Specification

PKT-SP-ESG-C01-170405

CLOSED

Notice

This PacketCable specification is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs® in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third-party documents, including open source licenses, if any.

© Cable Television Laboratories, Inc. 2010-2017

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs and its members shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	PKT-SP-ESG-C01-170405			
Document Title:	Enterprise SIP Gateway Specification			
Revision History:	I01 - Released November 3, 2010 C01 - Released April 5, 2017			
Date:	April 5, 2017			
Status:	~~Work in Progress~~	~~Draft~~	~~Issued~~	**Closed**
Distribution Restrictions:	~~Author Only~~	~~CL/Member~~	~~CL/ Member/ Vendor~~	**Public**

Key to Document Status Codes

Work in Progress	An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
Draft	A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
Issued	A generally public document that has undergone Member and Technology Supplier review, cross-vendor interoperability, and is for Certification testing if applicable. Issued Specifications are subject to the Engineering Change Process.
Closed	A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

CableLabs® is a registered trademark of Cable Television Laboratories, Inc. Other CableLabs marks are listed at http://www.cablelabs.com/specs/certification/trademarks. All other marks are the property of their respective owners.

Contents

Figures

Tables

1 INTRODUCTION[1]

This specification defines the requirements for the PacketCable 2.0 Enterprise SIP Gateway (ESG) device. The primary purpose of the ESG is to simplify and streamline the initial deployment and ongoing management of Business Voice services to enterprise customers. The ESG sits at the boundary between the Service Provider and Enterprise network, and serves as a demarcation point between these two networks. It normalizes the wide variety of SIP (Session Initiation Protocol) signaling protocols supported by currently deployed enterprise CPE (Customer Premises Equipment) equipment into a single well-defined interface that is compatible with the PacketCable network. It also provides enhanced fault detection and reporting capabilities that speed up the detection, isolation, and resolution of service-affecting failures. Finally, the ESG can act as a Gateway device for provisioning traffic between the Service Provider network and operator-owned and managed Enterprise CPE equipment.

The ESG comes in two flavors - an embedded version where the ESG device contains a DOCSIS® Cable Modem, and a stand-alone version where the ESG device is separate from and connected to the PacketCable access network via a standard Ethernet interface.

1.1 Scope

All the normative requirements for the ESG are contained in this single specification. It defines a reference architecture that identifies new reference points connecting the ESG to the PacketCable 2.0 network architecture defined in [PKT-ARCH-TR]. It defines the ESG signaling requirements to support these new reference points in order to support the following functions:

- Call Control
- Media
- Provisioning and Management
- Quality of Service
- Security

1.2 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.

[1] Updated by ESG-N-12.0691-8 on 11/4/16 by PO

"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

2 REFERENCES[2]

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references.

All references are subject to revision, and parties to agreement based on this specification are encouraged to investigate the possibility of applying the most recent editions of the documents listed below.

[CL-CANN-DHCP-REG]	CableLabs DHCP Options Registry Specification, CL-SP-CANN-DHCP-Reg-I14-170111, January 11, 2017, Cable Television Laboratories, Inc.
[CL-SP-MIB-BB]	CableLabs Battery Backup MIB Specification CL-SP-MIB-BB-I04-100608, June 8, 2010, Cable Television Laboratories, Inc.
[eDOCSIS]	eDOCSIS Specification, CM-SP-eDOCSIS-I28-150305, March 5, 2015, Cable Television Laboratories, Inc.
[RFC 6076]	IETF 6076, Basic Telephony SIP End-to-End Performance Metrics. D. Malas, A. Morton. January 2011.
[RFC 6035]	IETF 6035, Session Initiation Protocol Event Package for Voice Quality Reporting. A. Pendleton, A. Clark, A. Johnston, H. Sinnreich. November 2010.
[PKT 24.229]	PacketCable SIP and SDP Stage 3 Specification 3GPP TS 24.229, PKT-SP-24.229-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT1.5-SEC]	PacketCable 1.5 Security Specification, PKT-SP-SEC1.5-I03-090624, June 24, 2009, Cable Television Laboratories, Inc.
[PKT-BSSF]	PacketCable Business SIP Services Feature Specification, PKT-SP-BSSF-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-EUE-DATA]	PacketCable E-UE Provisioning Data Model Specification, PKT-SP-EUE-DATA-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-EUE-PROV]	PacketCable E-UE Provisioning Framework Specification, PKT-SP-EUE-PROV-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-RST-E-DVA]	PacketCable Residential SIP Telephony E-DVA Specification, PKT-SP-RST-E-DVA-C01-140314,-March 14, 2014, Cable Television Laboratories, Inc.
[PKT-RST-EUE-PROV]	PacketCable RST E-UE Provisioning Specification, PKT-SP-RST-EUE-PROV-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-UE-PROV]	PacketCable 2.0 UE Provisioning Framework, PKT-SP-UE-PROV-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.

[2] Updated by ESG-N-12.0682-8 10/28/16, ESG-N-11.0665-6, 11/1/16 and ESG-N-12.0691-8 on 11/4/16 by PO

[PKT-CODEC-MEDIA] PacketCable 2.0 CODEC and MEDIA Specification, PKT-SP-CODEC-MEDIA-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.

[RFC 2131] IETF RFC 2131, Dynamic Host Configuration Protocol, March 1997.

[RFC 2663] IETF RFC 2663, IP Network Address Translator (NAT) Terminology and Considerations, August 1999.

[RFC 3264] IETF RFC 3264, An Offer/Answer Model with Session Description Protocol (SDP), June 2002.

[RFC 3315] IETF RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), July 2003.

[RFC 3329] IETF RFC 3329, Security Mechanism Agreement for the Session Initiation Protocol, January 2003.

[RFC 3484] IETF RFC 3484, Default Address Selection for Internet Protocol version 6 (IPv6), February 2003.

[RFC 3550] IETF RFC 3550/STD0064, RTP: A Transport Protocol for Real-Time Applications, July 2003.

[RFC 3611] IETF RFC 3611, RTP Control Protocol Extended Reports (RTCP XR), November 2003.

[RFC 4787] IETF RFC 4787, Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, January 2007.

[RFC 5245] IETF RFC 5245, Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols, April 2010.

[RFC 5393] IETF RFC 5393 Addressing an Amplification Vulnerability in Session Initiation Protocol (SIP) Forking Proxies, December 2008.

[RFC 5761] IETF RFC 5761, Multiplexing RTP Data and Control Packets on a Single Port, April 2010.

[RSTF] PacketCable Residential SIP Telephony Feature Specification, PKT-SP-RSTF-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.

[SIPconnect1.1] SIP Forum SIPconnect 1.1 Technical Recommendation Final (v27), SIP Forum, March 2011.

[TR-069] Broadband Forum TR-069, CPE WAN Management Protocol, Issue 1 Amendment 4, July 2011.

2.2 Informative References

This specification uses the following informative references.

[DOCSIS RFIv2.0] DOCSIS Radio Frequency Interface Specification, DOCSIS CM-SP-RFIv2.0-C02-090422, April 22, 2009, Cable Television Laboratories, Inc.

[ISO/IEC 19501] ISO/IEC 19501:2005 Information technology - Open Distributed Processing – Unified Modeling Language (UML) Version 1.4.2.

[PCMM] PacketCable Multimedia Specification PKT-SP-MM-I07-151111, November 11, 2015, Cable Television Laboratories, Inc.

[PKT-ARCH-TR] PacketCable Architecture Framework Technical Report, PKT-TR-ARCH-FRM-C01-140314, March14, 2014, Cable Television Laboratories, Inc.

[PKT-PROV1.5] PacketCable MTA Device Provisioning Specification, PKT-SP-PROV1.5-I04-090624, June 24, 2009, Cable Television Laboratories, Inc.

[PKT-QoS] PacketCable 2.0 Quality of Service Specifications, PKT-SP-QOS-C01-140314, March14, 2014, Cable Television Laboratories, Inc.

[RFC 3261] IETF RFC 3261, SIP: Session Initiation Protocol, June 2002.

[RFC 3966] IETF RFC 3966 The tel URI for Telephone Numbers, December 2004.

[RFC 3551] IETF RFC 3551/STD0065. RTP Profile for Audio and Video Conferences with Minimal Control, July 2003.

[RFC 3986] IETF RFC 3986/STD0066, Uniform Resource Identifier (URI): Generic Syntax, January 2005.

[RFC 4566] IETF RFC 4566, SDP: Session Description Protocol, July 2006.

[RFC 5626] IETF RFC 5626, Managing Client-Initiated Connections in the Session Initiation Protocol (SIP), October 2009.

[PKT-TR-DS-IPv6] PacketCable Dual-Stack IPv6 Architecture Technical Report, PKT-TR-DS-IP6-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027;
  Phone +1-303-661-9100; Fax +1-303-661-9199; http://www.cablelabs.com
- Broadband Forum, http://www.broadband-forum.org/
- Internet Engineering Task Force (IETF) Secretariat, 48377 Fremont Blvd., Suite 117, Fremont, California 94538, USA, Phone: +1-510-492-4080, Fax: +1-510-492-4001, http://www.ietf.org
- SIP Forum, Internet: http://sipforum.com/

3 TERMS AND DEFINITIONS[3]

This specification uses the following terms:

Application Level Gateway	An Application Level Gateway (ALG) provides transparent routing of IP messages between two IP address realms when the application-payload of the messages contains address information that is local to one of the address realms. In the context of this document, the ALG provides transparent routing between the private address realm of the Enterprise network and the public address realm of the Service Provider network.
Back-to-Back User Agent	A back-to-back user agent (B2BUA) is a logical entity defined in [RFC 3261]. It receives a SIP request and processes it as a user agent server (UAS) up through the SIP protocol layers to the Transaction User (TU) layer, where it is passed via undefined application logic to the TU of a user agent client (UAC). The UAC then generates a request based on the received TU event. Responses received by the UAC are passed to the UAS in the reverse direction. The B2BUA is therefore a concatenation of a UAC and UAS. No explicit definition is defined for the application behavior.
Business Voice	The collection of voice services provided to an enterprise customer. Business Voice includes two deployment models; SIP Trunking Service, and Hosted IP Centrex service.
Configuration Server	The logical network element responsible for UE provisioning, configuration and management.
CWMP ALG	A CWMP ALG provides transparent routing of CWMP messages between the Service Provider network and the Enterprise network.
Dual-Stack mode	A configuration option where the ESG acquires and uses WAN IP addresses of both IP versions (IPv4 and IPv6).
Enterprise SIP Entity	A SIP-PBX or a SIP endpoint, located in the Enterprise network.
Gateway	A Gateway is an entity that enables communication between address realms (e.g., between the private address realm of the Enterprise network and the public address realm of the Service Provider network).
Hosted IP Centrex Service	The business voice deployment model where service control for enterprise users resides in the Service Provider network. This is referred to as "Business SIP Services" in PacketCable 2.0. The enterprise SIP entity is a SIP endpoint.
HTTP ALG	An HTTP ALG provides transparent routing of HTTP messages between two IP address realms when the application payload of the HTTP message contains address information that is local to one of the address reams.
Network Address Translator	A Network Address Translator (NAT) manipulates the IP address:port information in the IP and TCP/UDP headers to provide transparent routing of IP packets between two IP address realms (e.g., between the private address realm of the Enterprise network and the public address realm of the Service Provider network).
Provisioning Gateway	A Gateway that enables exchange of provisioning traffic between address realms.
Provisioning Server	See Configuration Server.

[3] Updated by ESG-N-12.0691-8 on 11/4/16 by PO

RTCP packet	A control packet consisting of a fixed header part similar to that of RTP data packets, followed by structured elements that vary depending upon the RTCP packet type. Typically, multiple RTCP packets are sent together as a compound RTCP packet in a single packet of the underlying protocol; this is enabled by the length field in the fixed header of each RTCP packet.
RTP packet	A data packet consisting of the fixed RTP header, a list of contributing sources, and the payload data.
SIP Endpoint	See Enterprise SIP Entity.
SIP-PBX	A Private Branch eXchange (PBX) deployed in the enterprise network, where the network-facing interface to the cable Service Provider is SIP. Business voice service control for the enterprise users resides at the SIP-PBX.
SIP phone	See Enterprise SIP Entity.
SIP Trunking Service	The Business Voice deployment model where the Service Provider network provides network connectivity to a SIP-PBX located in the Enterprise network.
Traditional NAT	A Traditional NAT updates the source IP address:port information for messages sent from the private to the public address realm, and updates the destination IP address:port for messages sent from the public to the private address realm.
Twice NAT	A Twice NAT updates both the source and the destination IP address:port information for messages exchanged between the two address realms.

4 ABBREVIATIONS AND ACRONYMS[4]

This specification uses the following abbreviations:

ACS	Auto-Configuration Server
ALG	Application Level Gateway
B2BUA	Back-to-Back User Agent
CM	Cable Modem
CPE	Customer Premises Equipment
CWMP	CPE WAN Management Protocol
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
E-ESG	Embedded Enterprise SIP Gateway
eSAFE	Embedded Service/Application Functional Entity
ESE	Enterprise SIP Entity
ESG	Enterprise SIP Gateway
GUA	Global Unicast Address
IBCF	Interconnection Border Control Function
ID	Identifier
IETF	Internet Engineering Task Force
IMPI	IP Multimedia Private Identity
IMPU	IP Multimedia Public Identity
IMS	IP Multimedia Subsystem
IP	Internet Protocol
LAN	Local Area Network
MIB	Management Information Base
NAPT	Network Address Port Translation
NAT	Network Address Translation
PC 2.0	PacketCable 2.0
P-CSCF	Proxy - Call Session Control Function
QoS	Quality of Service
RFC	Request For Comment
SBC	Session Border Controller
SDP	Session Description Protocol
S-ESG	Standalone Enterprise SIP Gateway
SETA	SIP Endpoint Test Agent

[4] Updated by ESG-N-12.0682-8 10/28/16, ESG-N-11.0665-6, 11/1/16 ESG-N-12.0691-8 on 11/4/16 by PO

SIP	Session Initiation Protocol
SP-SSE	Service Provider SIP Signaling Entity
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TLV	Type/Length/Value
UDP	User Datagram Protocol
UE	User Equipment
UML	Unified Modeling Language
VoIP	Voice over Internet Protocol
WAN	Wide Area Network

5 OVERVIEW[5]

Figure 1 shows a high-level view of the ESG, and its relationship and interconnectivity to the Service Provider and Enterprise networks in the delivery of business voice services to enterprise customers. The ESG provides a well-defined, managed demarcation point at the inter-network boundary that greatly simplifies the operator's task of detecting and isolating service-affecting problems. The ESG also provides a standard interface to the Service Provider network, thus enabling the Service Provider to more easily deploy business voice services to a wide variety of diverse voice CPE equipment.

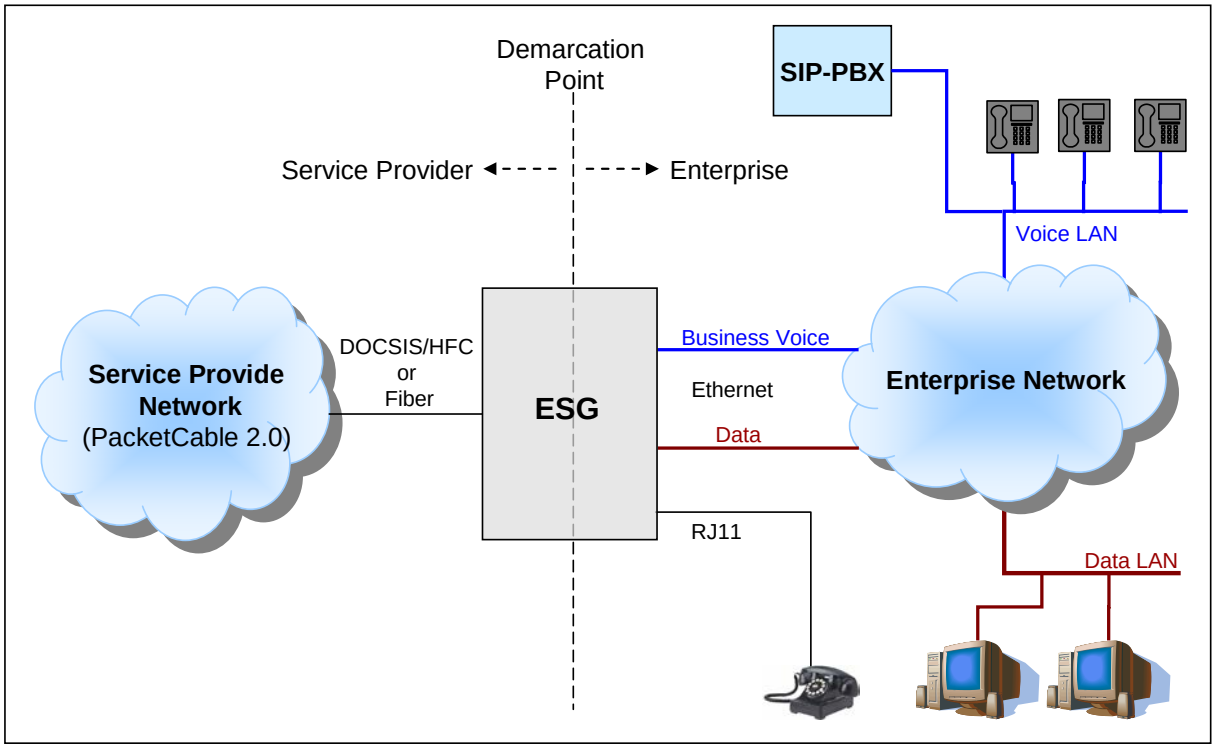

Figure 1 - ESG Network Architecture Overview

Within the overall scope of business voice services, the ESG must support a fairly wide variety of deployment and service scenarios. For example, this specification describes two versions of the product: one where the ESG is embedded in a DOCSIS Cable Modem (CM) and supports an HFC port toward the SP Network, and one where the ESG is a stand-alone device supporting a standard Ethernet port that can be connected to either DOCSIS/HFC or fiber transport to the SP network. The embedded version of the ESG also supports up to 4 E-DVA analog line RJ11 ports to support enterprise FAX and alarm panels, and a standard Ethernet RJ45 port to provide broadband data service to the enterprise.

Note: The term "DOCSIS" in this document is understood to refer to DOCSIS specification version 1.1 or later unless explicitly stated otherwise. Please refer to the corresponding DOCSIS specifications for more information about DOCSIS (for instance, DOCSIS 2.0 is specified in [DOCSIS RFIv2.0] and associated specifications).

[5] Updated by ESG-N-12.0691-8 on 11/4/16 by PO

Business Voice services include support for both SIP Trunking service to a SIP-PBX, and support of hosted IP Centrex service to enterprise SIP endpoints. Referring to Figure 1, the SIP procedures supported on the ESG interface to the PacketCable 2.0 network, and the point of connection at the SIP layer into the PacketCable 2.0 network depends on the service as follows:

1. For SIP Trunking service, the ESG supports SIP procedures toward the PacketCable 2.0 network that comply with the SIP-PBX procedures defined in SIPconnect 1.1 Technical Recommendation [SIPconnect1.1]. SIPconnect1.1 defines two modes of operation; the Registration Mode and the Static Mode.

    a) When operating in the Registration Mode, the SIP-PBX conveys its SIP signaling address to the PacketCable network dynamically, using a variant of the SIP registration procedure as defined in [SIPconnect1.1]. This procedure enables the SIP-PBX to register all the enterprise users with the PacketCable network using a single SIP registration transaction. The SIP signaling interface point into the PacketCable network is at the Proxy - Call Session Control Function (P-CSCF).

    b) When operating in the Static Mode, the PacketCable network views the SIP-PBX as a peer network, where the SIP signaling interface point into the PacketCable network is at the Interconnection Border Control Function (IBCF).

2. For hosted IP Centrex service, the ESG supports SIP procedures toward the PacketCable 2.0 network that comply with the DVA requirements defined in the PacketCable 2.0 Business SIP Services (BSS) Feature Specification [PKT-BSSF]. The PacketCable network views the enterprise users as BSS business users that register directly with the home PacketCable network. The SIP signaling interface point into the PacketCable network is at the P-CSCF. The ESG also supports a Provisioning Gateway function that provides transparent routing of provisioning messages between the hosted SIP endpoints in the Enterprise network and the Provisioning Server in the Service Provider network.

5.1 ESG Call Signaling and Control Functions

Figure 2 shows an example of the internal architecture of the call signaling and control functions within the ESG. This diagram is included for illustrative purposes only, as a means of describing the behavior of the ESG, and is not meant to mandate a specific implementation.

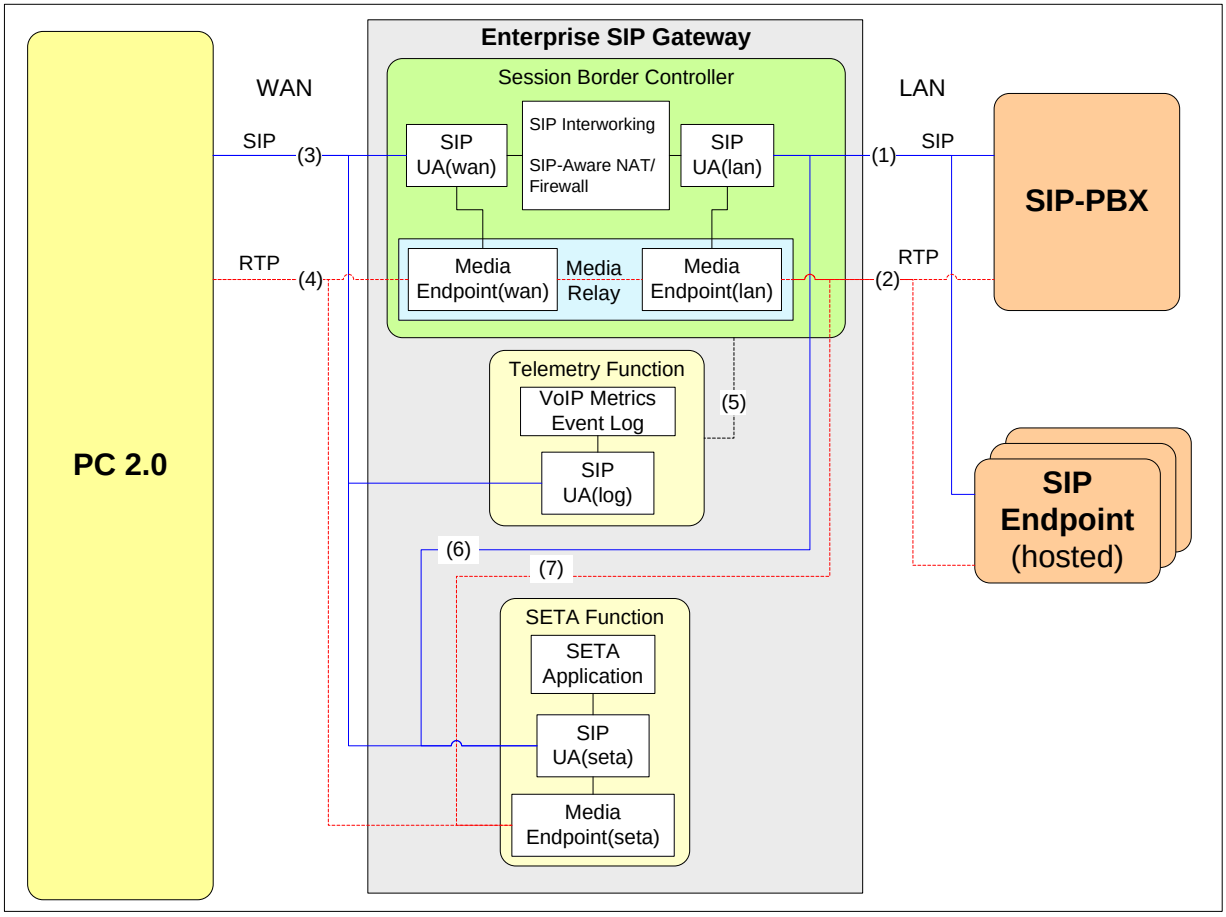

Figure 2 - ESG Block Diagram

The ESG connects the enterprise Local-Area-Network (LAN) to the Service Provider Wide-Area-Network (WAN) in order to carry business voice traffic between the enterprise SIP entities (SIP-PBX or hosted SIP endpoints) and the PacketCable 2.0 network. Starting on the right-hand-side of the diagram, interfaces (1) and (2) carry SIP signaling and RTP media respectively between the enterprise SIP entities and the LAN side of the ESG. Interfaces (3) and (4) in turn carry the SIP signaling and RTP media respectively between the WAN side of the ESG and the PacketCable 2.0 Network. The ESG contains a Session Border Controller (SBC), a SIP Endpoint Test Agent (SETA) Function, and a Telemetry Function. These functions monitor and manipulate the SIP messages and RTP packets as they pass between the LAN and WAN interfaces on the ESG.

5.1.1 Session Border Controller (aka SIP ALG)

The Session Border Controller (SBC) supports three separate functions: a SIP Interworking function, a SIP-aware Network Address Translation (NAT), and a SIP-aware firewall.

Note: The ESG Session Border Controller falls into the general category of functions commonly referred to as Application Level Gateways (ALG). In this case the SBC is a ***SIP ALG***, to reflect the fact that the ALG has knowledge of SIP and its related VoIP protocols SDP, RTP, and RTCP.

The NAT function is mandatory to implement and use; i.e., the ESG is located at the demarcation point between the Service Provider and Enterprise network, and as a SIP-aware NAT it manipulates IP addresses in the IP header, SIP headers, SIP body (SDP), and RTP/RTCP headers to translate between the LAN-side Enterprise IP addresses and the

WAN-side Service Provider IP addresses. The NAT function also supports IPv4-to-6 version interworking, say when a Service Provider network supporting IPv6 wants to provide service to an Enterprise network that supports IPv4.

The SIP firewall and Signal Normalization functions are mandatory to implement. However, these functions are optional to use in the sense that the firewall rules and interworking procedures are configured by the operator, and can be configured to "no firewall rules" and "no interworking procedures" which effectively disables these functions. For example, if the SIP-PBX is fully compatible with the Service Provider network, then the operator can choose to configure the SBC such that it applies no interworking procedures, and hence (aside from the NAT function) becomes transparent to SIP signaling exchanged between the Enterprise and the Service Provider network.

Figure 2 shows an implementation example where the SBC is implemented as a Back-to-Back-User-Agent (B2BUA). SIP UA(wan) faces the Service Provider network, and supports SIP procedures compatible with PacketCable 2.0 on interface (3), while SIP UA(lan) faces the enterprise network, and supports the SIP procedures compatible with the SIP-PBX or hosted SIP endpoint on interface (1). These back-to-back SIP UAs are connected by an Interworking Function that applies SIP header manipulation rules plus any other signal normalization procedures required to achieve interworking between interfaces (1) and (3). It also enforces the SIP firewall rules and performs the NAT functions.

5.1.2 Telemetry Function

The purpose of the Telemetry Function shown in Figure 2 is to collect data as an aid in detecting and resolving problems. The Telemetry Function collects the following data from the SBC over interface (5):

- VoIP Metrics (Voice over Internet Protocol)
- Error events such as 4xx, 5xx and 6xx responses to SIP INVITE
- SIP and RTP message traces.

Note: Interface (5) is internal to the ESG, and therefore, not defined in this specification.

The Telemetry Function can report this data autonomously to the SP network (e.g., report VoIP Metrics), can upload the data to the SP network based on operator request (e.g., upload trace files), and can report alarms associated with the data (e.g., alarm indicating poor voice quality). The SIP UA(log) reports VoIP Metrics to the SP network in a PUBLISH request over interface (3).

5.1.3 SETA Function

The purpose of the SETA Function is to initiate and accept test calls under management control in order to verify the health of the ESG and its connectivity to the SP Network. The operator can use the management interface to initiate test calls on demand, or at a programmed periodic interval. The SETA can also accept test calls, including RTP loopback calls (RTP packet reflector only).

As shown in Figure 2, the SETA can exchange its SIP messages and RTP packets directly with the PacketCable 2.0 network via interfaces (3) and (4). SETA can also be configured to exchange SIP and RTP with the LAN side of the SBC, via interfaces (6) and (7). This configuration option causes the SETA SIP and RTP traffic to traverse the SBC on their way to the PacketCable 2.0 network, thus enabling SETA to verify basic SBC functionality. Since interfaces (6) and (7) are internal, the details of how this is accomplished are not specified in this document.

5.1.4 ESG Call Signaling/Control Function Deployment Options

Figure 3 shows a deployment scenario where the embedded version of the ESG is serving an enterprise network that separates the voice and high-speed data on separate physical LANs. The ESG supports two Ethernet ports toward the enterprise: one dedicated for voice, and one dedicated for data. The ESG data port is connected to the enterprise router/firewall. The ESG voice port is connected directly to the enterprise voice LAN, effectively bypassing the

enterprise router/firewall. In this case the ESG-SBC serves as the firewall (a SIP-aware firewall) for the voice traffic between the enterprise and the outside world.

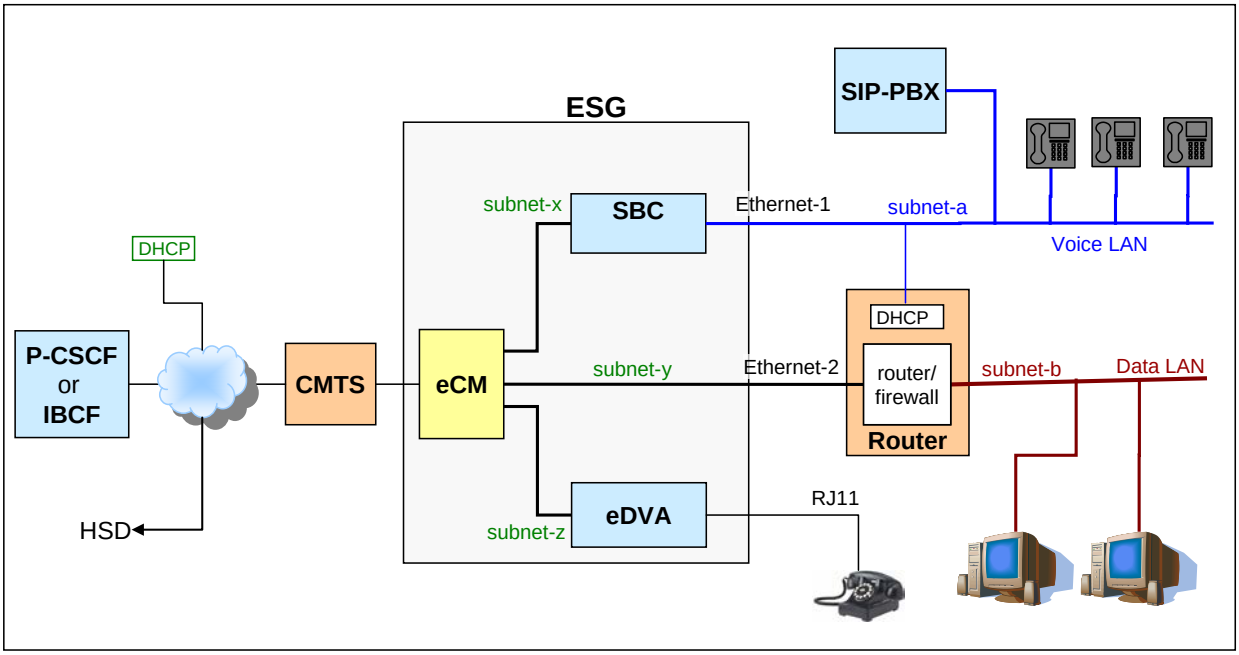

Figure 3 - ESG Deployment Option - Voice & Data on Separate LANs

As shown in Figure 3, the SBC obtains a public IP address from the Service Provider Dynamic Host Configuration Protocol (DHCP) server (in this example, an IP address on subnet-x). Likewise, the WAN side of the enterprise Router and the eDVA obtain IP addresses from the SP DHCP server on subnet-y and subnet-z respectively. The SBC obtains a private IP address from the enterprise DHCP server, in this example on subnet-a, which is the same subnet used by the SIP-PBX.

Separating the voice and data traffic onto separate physical LANs ensures that data traffic will not adversely affect the voice service. In deployments where the enterprise network combines voice and data traffic on the same LAN, other techniques such as VLAN tagging can be used to provide adequate quality-of-service for the voice traffic.

The ESG will always separate physical voice and data interfaces to the enterprise network. In deployments where the voice and data are in fact sharing the same LAN within the enterprise, the method of ensuring that data doesn't interfere with voice traffic is out-of-scope of this specification.

Note: A future release of this specification may define a single-port version of the ESG, where the ESG adds VLAN tagging capability and combines voice and data on a single Ethernet port toward the enterprise.

5.1.5 ESG Call Signaling/Control Function Assumptions

The ESG architecture and requirements in this document are based on the following key assumptions:

- ESG access to the Service Provider (SP) is via network facilities controlled and managed by the Service Provider. Issues related to access over a third-party unmanaged network facility are not addressed by this specification.
- The ESG and SP network are directly connected without any NAPT (Network Address Port Translation) device between them.
- The ESG and the Enterprise SIP entity (ESE) are directly connected without any NAPT device between them.

- The ESG receives its WAN-side IP address from the Service Provider using mechanisms such as DHCP or manual provisioning.
- The ESG receives its LAN-side IP address from the Enterprise using mechanisms such as DHCP or manual provisioning. Network and device configuration to achieve this are out of scope of this specification.
- The ESE (e.g., SIP-PBX) is configured to communicate with the PacketCable 2.0 network via the LAN-side IP address of the ESG.
- The ESG provides two LAN-side Ethernet ports; a "Voice" port dedicated for voice services, and a "Data" port dedicated for high-speed data.
- Quality of Service (QoS) in the Enterprise network (e.g., between the ESE and ESG) is out of scope of this specification.

5.2 Provisioning Gateway[6]

When deployed for hosted IP-Centrex service, the ESG can serve as a Provisioning Gateway at the demarcation point for provisioning traffic exchanged between the Service Provider's provisioning servers and the enterprise SIP phones. Support of this function in the ESG provides operator-controlled access to the provisioning interface on the SIP endpoints and resolves the problem where the enterprise NAT/Firewall is configured to block this provisioning traffic.

The remainder of Section 5.2 provides an informative overview of the different types of Provisioning Gateways supported by the ESG. Section 6.6 describes the normative requirements for a Provisioning Gateway that is acting as a Data NAT/Firewall. Section 8.3 defines the normative requirements for a Provisioning Gateway that is acting as a Provisioning ALG. Annex A.1.4 provides an informative description of the provisioning message flows supported by the Provisioning Gateway. Annex A.2 defines the data model for the Provisioning Gateway.

5.2.1 Types of Gateways

Depending on the deployment scenario, the Provisioning Gateway acts either as a Network Address Translator (NAT), or a provisioning-aware Application Level Gateway (ALG). Two types of NATs are described; the Traditional NAT and the Twice NAT. Likewise, two types of Provisioning ALG's are described; the HTTP ALG and the CWMP ALG.

Figure 4 shows a "Class Hierarchy" view of the different types of Provisioning Gateways supported by the ESG.

[6] New sections added by ESG-N-12.0691-8 on 11/4/16 by PO

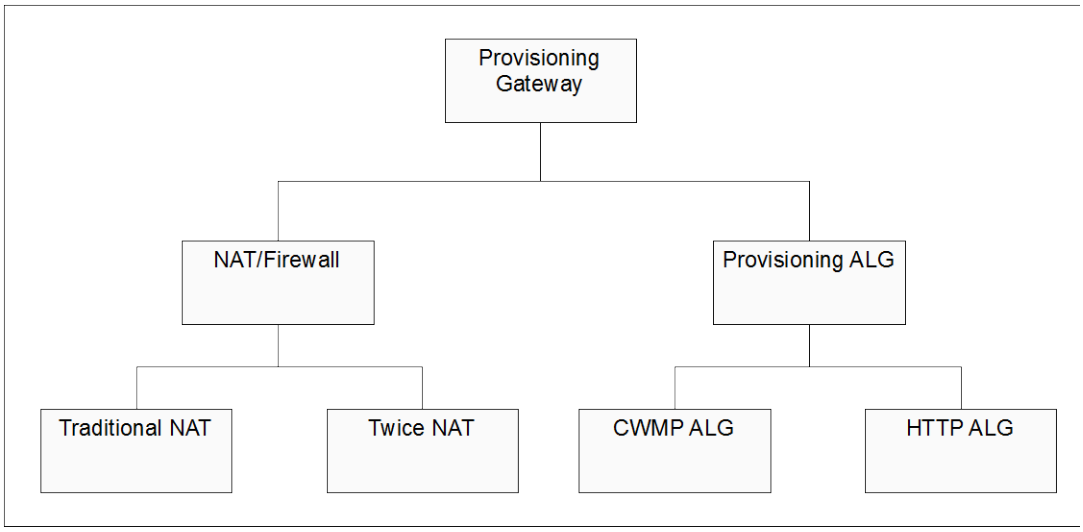

Figure 4 - Provisioning Gateway Hierarchy Diagram

An MSO would choose to deploy the type of Provisioning Gateway that best suits their specific hosted IP-Centrex service deployment scenario.

5.2.1.1 Network Address Translator (NAT)

A NAT enables transparent routing of IP messages between hosts in two different address realms. (In the context of the ESG, the two address realms are the private Enterprise LAN network and the public Service Provider WAN network.) The NAT achieves this by modifying the IP address:port information carried in the IP and TCP/UDP headers for messages exchanged between the two hosts. The NAT also maintains session state information in the form of address mappings, so that messages within a TCP connection, or UDP messages within a request/response transaction, can be delivered to the correct host.

[RFC 2663] describes two types of NATs that apply to the Provisioning Gateway function; the Traditional NAT and the Twice NAT. Both of these NAT-types support unidirectional sessions, which are outbound sessions initiated from the private LAN network to the public WAN network. The difference between these NATs is that the Traditional NAT updates either the source address or destination address (not both) for datagrams crossing address realms, while the Twice NAT updates both the source and destination addresses as datagrams cross address realms.

Traditional NAT: The SIP phone places the public IP address of the provisioning server in the destination address of the provisioning request, and its own private IP address as the source address, before sending the request to the LAN interface of the NAT. On receiving the request from the phone, the NAT knows that it can't send it to the provisioning server as-is, because the source address (which is used to route the response to the request) isn't publically routable. Therefore, the NAT assigns a temporary public IP address:port on the WAN interface of the NAT, and creates a session that binds this NAT WAN IP address:port to the phone's LAN IP address:port. The NAT then updates the source IP address of the request with the newly assigned WAN public IP address:port, and forwards the request on to the provisioning server. When the NAT receives a subsequent response to the request on its WAN IP address:port of the session, it updates the destination address of the response to the mapped LAN IP address:port and forwards the response on to the phone.

Twice NAT: The Twice NAT works the same as a traditional NAT, except that the SIP phone doesn't know the public IP address of the provisioning server. Instead, the phone is configured with routing information so that it sends provisioning requests to a specific private LAN IP address:port on the NAT. The NAT is configured with address mapping information that maps this LAN IP address:port to the public IP address:port of the target server. Therefore, in addition to the address mapping duties of the Traditional NAT, the Twice NAT updates the

destination address of requests from the phone to the provisioning server, and source address of responses from the provisioning server to the phone.

The NAT has no application-level knowledge. Therefore, the NAT does not perform LAN/WAN interworking of address information carried within the application payload portion of the IP packet, nor can it establish additional address-mapping sessions that may be signaled in the application payload. If the application requires these additional application-specific routing functions, then additional mechanisms are required above and beyond those provided by NAT; either the ALG function described in Section 5.2.1.2, or a client-initiated application-level NAT traversal mechanism such as the STUN-based mechanism described in [TR-069].

5.2.1.2 Provisioning ALG

ALG functionality is required when IP address:port interworking within the IP and TCP/UDP headers isn't enough; when the Gateway device requires additional application-level knowledge to enable transparent routing of messages between hosts in two different address realms. In addition to NAT WAN/LAN interworking at the IP and TCP/UDP layers, the ALG provides two functions:

- It updates application payload information (usually address information) that is specific to the LAN or WAN networks as messages cross address realms.
- It establishes additional NAT sessions, where a session is WAN-to-LAN address binding, to enable subsequent packets associated with that session to be routed across address realms. For example, a CWMP-aware ALG establishes a session based on the CWMP Inform ConnectionRequestURL so it can route subsequent unsolicited connection requests from the ACS to the SIP phone.

The ALG is configured with a listening LAN IP address:port, and the publicly routable address of the server associated with that LAN address:port. Therefore, when the ALG receives a request on the configured LAN listening port, it forwards the request on to the configured server. In addition to updating the destination the IP address of the packet to identify the public address of the server, ALG performs the traditional NAT functionality described above - it assigns a source WAN IP address:port for the request and maintains address binding information to forward subsequent responses back to the SIP phone.

The address of the target server is typically carried in the application portion of the request (e.g., an HTTP request Request-URI contains the server's URL). Therefore, the ALG must reach in and update this application-related information before forwarding the received request on to the server. This application-level activity is what distinguishes the ALG from the traditional NAT described above. And since the ALG has application knowledge, it can update any other LAN address information in the application payload of the request, such as the LAN address of the phone. Therefore, by interworking the LAN/WAN address information in the application payload of the message, the ALG provides a NAT-traversal mechanism for a specific application protocol. (This specification describes two types of Provisioning ALGs; the CWMP ALG and the HTTP ALG).

5.2.2 Whether to Deploy Traditional NAT, Twice NAT, or Provisioning ALG

One factor that governs the operator's decision to deploy the ESG Provisioning Gateway as a Traditional NAT, Twice NAT, or an ALG, is whether the SIP phone knows the public address of its provisioning server. As a rule-of-thumb, if the phone knows the public IP address of the provisioning server, then the Provisioning Gateway can act as a Traditional NAT. If the phone doesn't know the public IP address of the provisioning server, and instead relies on the Provisioning Gateway to retarget provisioning requests to the public address of the server, then either a Twice NAT or an ALG are required.

Why not *always* provide the phone with the public IP address of the provisioning server? One reason is that in order to enable IP routing of provisioning traffic through the ESG, certain deployment arrangements require the phone to be configured with a provisioning server address that points to the ESG LAN interface. There may also be topology-hiding or provisioning process reasons for not pushing the Provisioning server address information all the way down to the enterprise phones.

5.2.2.1 Single vs. Multiple Gateways

An Enterprise network consisting of multiple LANs and subnets across a layer-2/layer-3 routing infrastructure typically applies the following high-level routing rules:

1) Packets destined for private IP addresses within the Enterprise network address space are routed directly to the enterprise host device that is assigned that address.

2) All other packets - packets addressed to public IP addresses outside the Enterprise address space - are routed to the same place; to the "default" Gateway that relays the packets on to the external network.

If there is only one Gateway device relaying requests to the Service Provider network, then that Gateway is designated as the default Gateway. In this case, the provisioning requests can be targeted to the provisioning server's pubic IP address and routed via the default Gateway to the external network, along with all the other data requests.

However, if there are two Gateway devices (say, the ESG for provisioning traffic, and an Enterprise NAT/Firewall for other data traffic), then only one of them (usually the Enterprise NAT/FW) is designated as the default Gateway. In this case, provisioning requests initiated by the enterprise SIP endpoint cannot be targeted to the public IP address of the provisioning server - since doing so would cause them to fall through to the default Gateway route and be sent to the Enterprise NAT/FW. Instead, the provisioning requests must be targeted to the private IP address of the ESG so that they are delivered to the LAN interface of the ESG.

This single vs. multiple Gateway routing consideration is illustrated in the following two examples.

In Figure 5 the ESG serves as the single gateway device at the demarcation point between the Service Provider and Enterprise network. The VoIP traffic (SIP, SDP, RTP, etc.) traverses the SIP ALG as described in Section 5.1.1. Provisioning traffic, along with the other high-speed data traffic, traverses the ESG NAT function. Since "Prov" and "Other Data" are both being routed through the same gateway (the ESG), the SIP phone can set the destination IP address of the provisioning requests to the public address of the provisioning server, so that provisioning requests fall through to the phone's default gateway route and are routed via the ESG along with the other data traffic. The ESG Gateway function for provisioning (and other) traffic can operate as a Traditional NAT, since it doesn't have to retarget the incoming requests to a new destination.

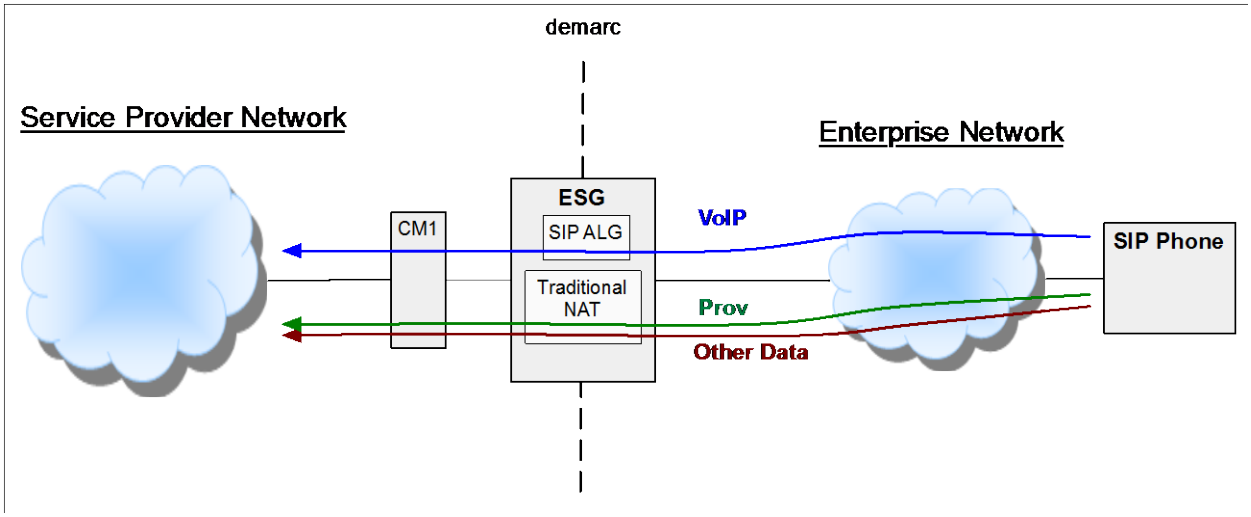

Figure 5 - ESG Traditional NAT/Firewall

Figure 6 shows the more common case where ESG shares gateway duties with an Enterprise-owned and operated NAT/Firewall. The clients in the Enterprise network (including the SIP phones) are configured with a default gateway route that points to the Enterprise NAT/Firewall. Since we want the provisioning requests from the SIP phone to be routed via the ESG, they cannot be sent to the public address of the provisioning server; otherwise they would use the phone's default gateway route and be routed via the Enterprise NAT/Firewall. In order to enable IP routing of provisioning requests to the ESG, the SIP phones are required to send provisioning requests to an ESG LAN IP address:port (see Section 5.2.3.1.2 for more details on IP routing). The Provisioning Gateway function in the ESG then retargets the requests received on its LAN address:port to the public address of the provisioning server. As shown in Figure 6, the Provisioning Gateway can perform this retargeting function in one of two ways:

- As a Twice NAT, where it updates the destination address in the IP header of incoming requests from the SIP phone, but does not modify any of the application payload information in the request, or
- As a Provisioning ALG, where it updates both the destination address in the IP header and the HTTP Request-URI to the public IP address and URI of the provisioning server. Since the Request-URI is part of the application payload, this gateway function is by definition an Application Level Gateway - in this case a Provisioning ALG. As a Provisioning ALG, it is able to update other address information carried in the application payload of provisioning messages.

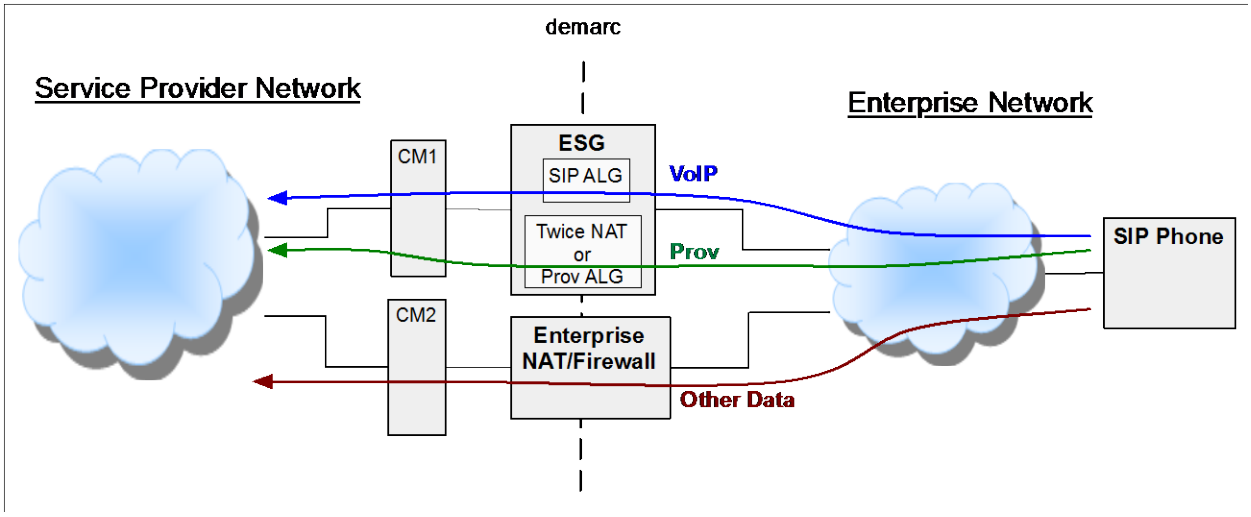

Figure 6 - ESG Twice-NAT or Provisioning ALG

5.2.2.2 Other NAT vs. ALG Considerations

In addition to the Enterprise network topology (single vs. multiple Gateways), other factors that influence whether the Provisioning Gateway is deployed as a NAT or an ALG include:

- **NAT-traversal:** if the provisioning protocol carries IP address information in its application payload, and the SIP phones and provisioning server don't support a suitable NAT-traversal mechanism (e.g., the STUN-based mechanism defined in [TR-069], then the ALG may be the right choice since it can perform the LAN/WAN address interworking of this application data on behalf of the phones and server.
- **Firewalls:** Most NATs are coupled with a Firewall that applies rules to allow or block messages exchanged between the Enterprise and Service Provider network. If the provisioning protocol supports special procedures that are blocked by a non-application-aware Firewall, then an ALG with its application-level knowledge may be required.
- **Security:** Since an ALG modifies the application payload of the provisioning messages, TLS sessions established by the SIP endpoint for provisioning messages must be hop-by-hop; from the SIP endpoint to the Provisioning ALG, and from the Provisioning ALG to the provisioning server. Therefore, if the provisioning

security model requires an end-to-end TLS connection between the SIP endpoint and the provisioning server, then the Provisioning Gateway must be deployed as a NAT.

5.2.3 Provisioning ALG Types (CWMP, HTTP)

The behavior of a Provisioning ALG must be tailored to support a specific provisioning protocol. The long-term goal is to adopt TR-069 as the standard provisioning model for Enterprise SIP Endpoints. Therefore, this document will focus on the requirements for an ALG that provides gateway services for the CWMP provisioning protocol defined by TR-069; a CWMP ALG.

Although the industry is moving toward adoption of TR-069 as a standard provisioning model, many of today's SIP phones support vendor-proprietary provisioning protocols. Although these provisioning protocols are non-standard, they are often based on HTTP. Therefore, to accommodate currently deployed SIP phones that don't yet support TR-069, this document will also define a general HTTP ALG function.

Note: The Provisioning ALG provides seamless interworking between the Enterprise LAN and Service Provider WAN networks when the SIP endpoints and the Service Provider provisioning servers support the same provisioning model. It is not expected to provide provisioning-interworking functions to enable support for the case where the SIP phones and network servers support different provisioning models/protocols.

5.2.3.1 CWMP ALG

The CWMP ALG provides interworking of CWMP provisioning messages between an Enterprise SIP Entity (ESE) located in the Enterprise network and an ACS located in the Service Provider network. The CWMP ALG is designed to support the specific case where the ESE is configured with an ACS URL that points to the LAN interface on the ESG; i.e., from the ESE's perspective, the ESG is its ACS. The ESG does not play the role of an ACS however. Rather, it supports a CWMP-aware gateway that provides WAN-LAN interworking at both the IP layer and application layer for CWMP messages exchanged between the ESE and the ACS located in the Service Provider network.

5.2.3.1.1 CWMP ALG Functions

In addition to performing traditional NAT functions for LAN/WAN IP address interworking at the IP and TCP/UDP layers, the CWMP ALG performs the following CWMP application-level functions:

- For CWMP requests received from the CWMP CPE (e.g., ESE):
  - Updates the destination IP address and HTTP Request URI to public address of the ACS
  - Provides LAN/WAN interworking of the ConnectionRequestURL received in the CWMP Inform request
  - Maintains WAN-to-LAN address mappings of the ConnectionRequestURL
  - Provides LAN/WAN interworking of the redirected URL received in a response from the ACS
- For CWMP connection requests received from the ACS:
  - Routes valid requests to the target ESE based on the ConnectionRequestURL address mappings,
  - Blocks invalid requests
- Supports TLS on both the LAN and WAN interfaces.

5.2.3.1.2 CWMP ALG Deployment Examples

This section describes how the CWMP ALG is integrated within the overall business services architecture, and how IP routing works in the Enterprise network for some common deployment examples.

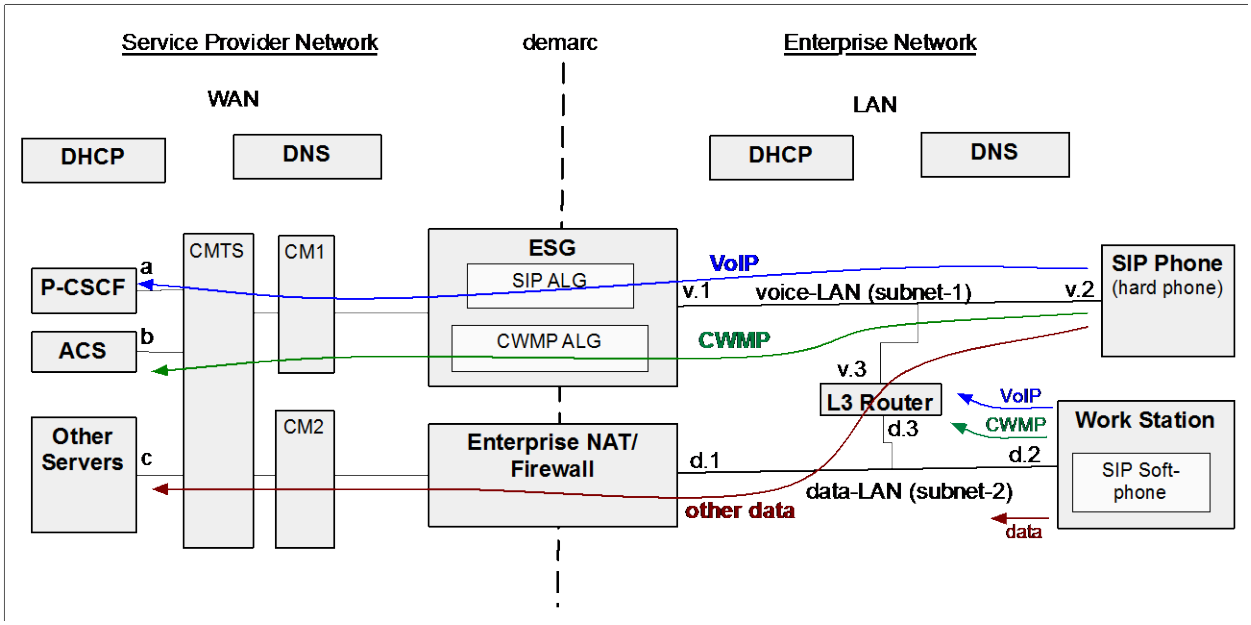

Figure 7 - CWMP ALG Deployment Example

Figure 7 shows the CWMP ALG within a business services deployment example where the Service Provider is providing both high-speed data and hosted IP-Centrex service to the enterprise customer.

The Service Provider network contains various servers, including:

- the P-CSCF, which provides VoIP signaling access to the PacketCable 2.0 network,
- the ACS, which is the TR-069-defined configuration server,
- other servers accessed by the enterprise users over the high-speed data service.

These network servers are assigned publicly routable IP addresses "a", "b", and "c", as shown in Figure 7.

The Enterprise network in this example contains two separate LANs; a voice LAN and a data LAN. The two LANs are inter-connected by a layer-3 router. The client devices in the Enterprise network include stand-alone SIP phones located on the voice LAN, and soft SIP phones running on general-purpose workstations located on the data LAN. The various hosts in the Enterprise network are assigned IP addresses "v.x" and "d.x" as shown in the diagram.

In this (typical) case the Service Provider deploys two Cable Modems (CMs); CM1 carries voice and provisioning traffic for the hosted SIP phones, while CM2 carries the remaining high-speed data traffic. The ESG and the Enterprise NAT/Firewall both sit on the demarcation point between the Service Provider and Enterprise network, providing NAT-traversal and access-control functions for all traffic between the Service Provider and Enterprise networks. The Enterprise owns the Enterprise NAT/Firewall, and therefore controls the firewall rule-set for high-speed data to/from the Service Provider network. The Service Provider owns the ESG, and therefore controls the NAT traversal and access control rules for the traffic that supports the hosted voice service; namely the VoIP protocols (SIP, SDP, RTP, RTCP) via the SIP ALG, and CWMP via the CWMP ALG.

In order to enable the different traffic types to be properly routed within the Enterprise network, the SIP Phones and workstations must be configured with the following IP routing information:

1) Routing info for SIP phones on the "voice" LAN:

- my P-CSCF = v.1:port-x

- my ACS = v.1:port-y
- default Gateway = v.3

The SIP phone is configured with the ESG LAN v.1 address for "my P-CSCF" and "my ACS". Since these are on the same subnet as the phone, the phone sends SIP and CWMP messages directly to ESG LAN interface. The ESG knows the public address of the target network P-CSCF or ACS, and updates the IP destination address to forward the request on to the correct server. For other (non-voice/CWMP) data traffic, the SIP Phone resolves the FQDN of the "other server" to a public IP address, and routes via its default-gateway route through the L3 Router, Enterprise NAT/FW, and CM2 to the Service Provider network.

2) Routing info for workstations on the "data" LAN that also support soft phones:

- my P-CSCF = v.1:port-x
- my ACS = v.1:port-y
- default Gateway = d.1

The soft-phone is configured with the same ACS and P-CSCF addresses as the SIP phone. Therefore, it sends upstream SIP requests to destination IP address v.1. Since the destination IP is not within the local subnet, the phone sends the message to its default gateway, which is the Enterprise NAT/Firewall LAN interface d.1. The Enterprise NAT/Firewall contains routing logic to forward the packet to the L3 router via d.3, which forwards the message on to the ESG LAN interface v.1. Upstream provisioning requests follow a similar route. Other data traffic (requests destined for public IPs) would be sent via default gateway routing to the Enterprise NAT/FW, CM2, and on to the Service Provider network.

Other Enterprise network would generally follow the above routing logic. Here are some common variations:

- L3 router and Enterprise NAT/Firewall integrated in a single device

SIP phone routing information would be as described above. The embedded L3 router would forward packets received on the data LAN and destined for v.1 to the ESG LAN interface.

- Voice LAN consists of multiple subnets

SIP phone routing information for P-CSCF and ACS would be as described above. SIP phones on a different subnet than the ESG LAN interface would be given a default gateway that identifies the L3 router leg connected to the phone's local subnet. The L3 routing rules would route packets destined for v.1 to the ESG LAN interface, and packets destined for public IP addresses to the Enterprise NAT/Firewall.

- ESG is located behind the Enterprise NAT/Firewall

The Enterprise NAT Firewall must be configured with a L3 tunnel that allows the ESG WAN interface to obtain a public IP address from the Service Provider network. The IP routing information in the SIP phones located in the Enterprise network is configured as described above.

5.2.3.2 HTTP ALG

The HTTP ALG supports a small subset of the requirements of the CWMP ALG. Like the CWMP ALG, it supports traditional NAT functions for LAN/WAN IP address interworking at the IP and TCP/UDP layers. In addition, the HTTP ALG updates the destination IP address and the Request URI of incoming HTTP requests to the publicly routable address of the provisioning server.

6 ENTERPRISE SIP GATEWAY APPLICATION REQUIREMENTS

This section contains technical application-level requirements common to both embedded and standalone versions of the ESG. The requirements specific to each version of the ESG are defined in Sections 7 and 8.

6.1 ESG Application Reference Architecture

Figure 8 identifies the reference points and components involved in carrying SIP signaling between the ESG, and the Service Provider and Enterprise networks.

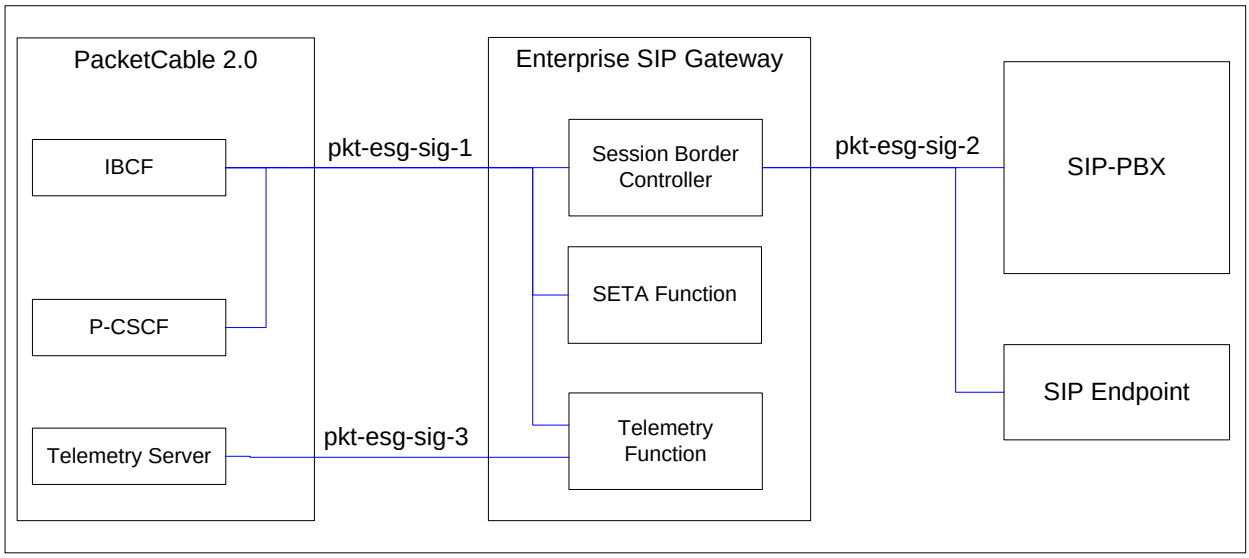

Figure 8 - ESG Signaling Reference Architecture

The reference points in Figure 8 are described in Table 1.

Table 1 - Signaling Reference Points Descriptions

Reference Point	PacketCable Network Components	Reference Point Description
pkt-esg-sig-1	IBCF - ESG/SBC IBCF - ESG/SETA P-CSCF - ESG/SBC P-CSCF - ESG/SETA	Carries SIP signaling between the IBCF and the ESG when the PacketCable 2.0 network is configured to interoperate with a SIP-PBX using the "static-mode" as defined in [SIPconnect1.1] (i.e., where the PC 2.0 network treats the SBC and the SETA as peer networks). Carries SIP signaling between the P-CSCF and the ESG when the PacketCable 2.0 network is configured to interoperate with a SIP-PBX using the "registration-mode" as defined in [SIPconnect1.1] (i.e., where the PC 2.0 network treats the SIP-PBX and the SETA as registered users).
pkt-esg-sig-2	ESG/SBC - SIP-PBX ESG/SBC - SIP Endpoint	Carries SIP signaling between the ESG SBC and a SIP-PBX or a SIP endpoint.
pkt-esg-sig-3	P-CSCF - ESG/Telemetry	The ESG Telemetry Function uses this interface to PUBLISH events to the Telemetry Collector.

Figure 9 identifies the reference points and components involved in carrying media (RTP & RTCP) packets between the ESG, and the Service Provider and Enterprise networks.

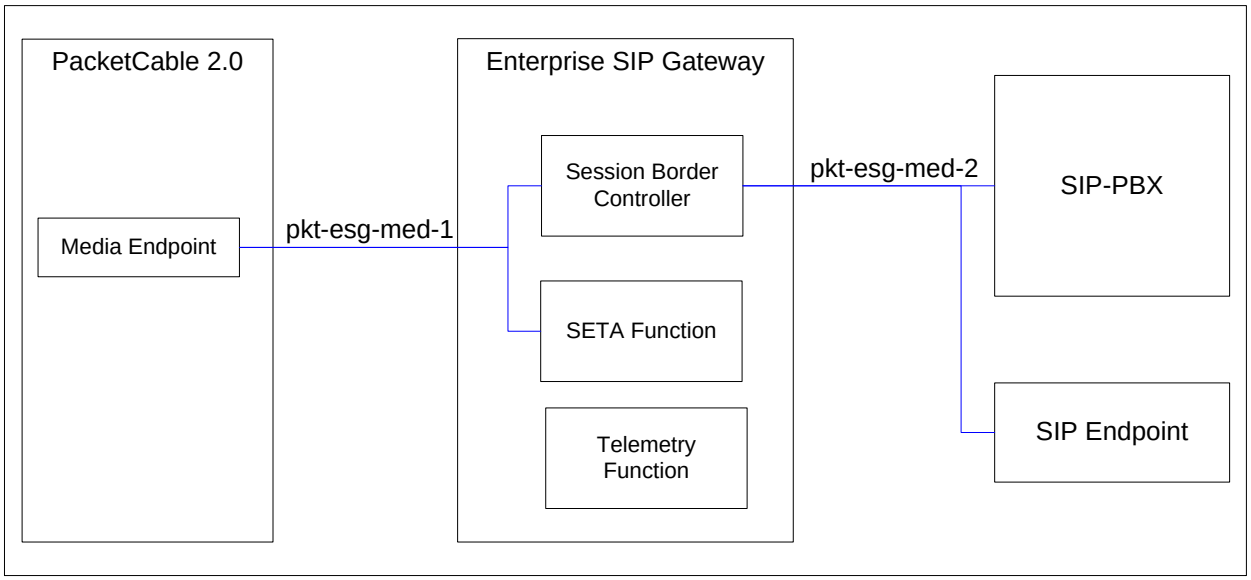

Figure 9 - ESG Media Reference Architecture

The reference points in Figure 9 are described in Table 2.

Table 2 - Media Reference Point Descriptions

Reference Point	PacketCable Network Components	Reference Point Description
pkt-esg-med-1	Media Endpoint - ESG/SBC Media Endpoint - ESG/SETA	Carries RTP/RTCP between a PacketCable 2.0 Media Endpoint (e.g., an E-DVA, Media Gateway, Media Server), and the WAN-side of the ESG.
pkt-esg-med-2	ESG/SBC - SIP-PBX ESG/SBC - SIP Endpoint	Carries RTP/RTCP between the LAN-side of the ESG and the SIP-PBX or hosted SIP Endpoint.

6.2 Quality of Service

This section defines the ESG requirements and functionality needed to support Quality of Service (QoS) on the PacketCable 2.0 access network.

6.2.1 Scope

This version of the specification assumes HFC using DOCSIS network protocol for the access network. QoS requirements for other access networks, such as DPoE™, are out of scope for this version of the specification. Additionally, as shown in Figure 10, the QoS procedures for the voice traffic between the Enterprise SIP Entity (e.g., SIP-PBX) and ESG, and between the S-ESG and Cable Modem (CM) are out of scope of this specification.

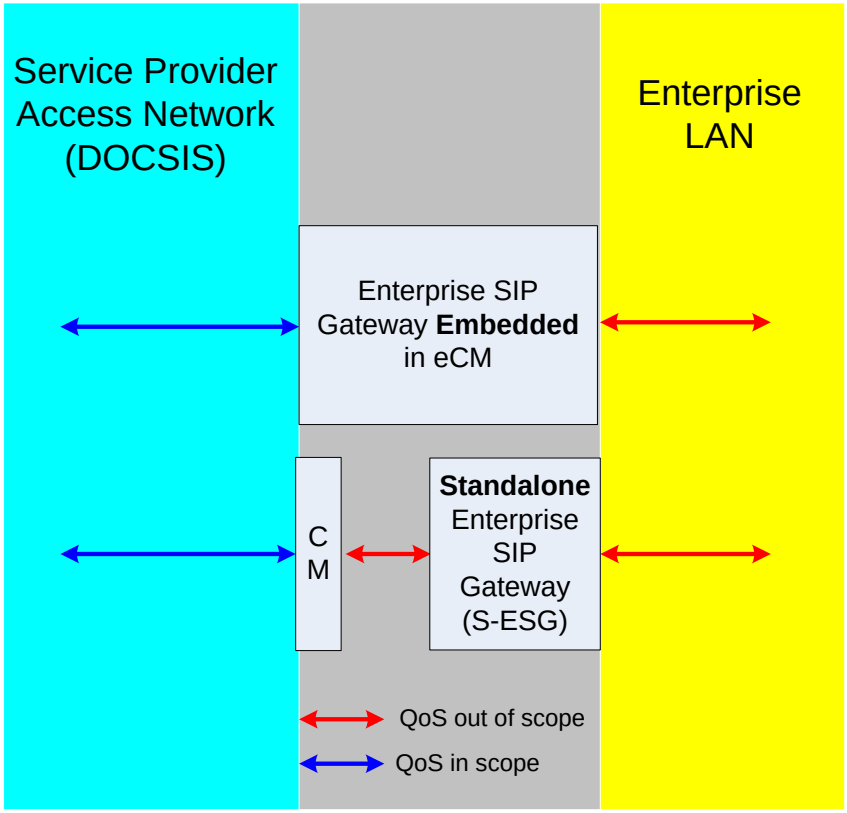

Figure 10 - Scope of QoS

6.2.2 Requirements

In PacketCable 2.0, QoS (Quality of Service) is handled using the PacketCable Multimedia (PCMM) mechanism. The ESG is intended to be QoS-unaware in this architecture, and therefore does not play a significant role in the allocation of QoS. However, the ESG does have some responsibilities in providing good voice quality to the Enterprise, and this section provides an overview of these QoS-related functions and the associated ESG requirements.

The PacketCable 2.0 core examines the SIP offer/answer and dynamically calculates the packet interval (usually 20 millisecond RTP frames), frame size, and classifier values. It then uses the interfaces described in [PKT-QoS] and [PCMM] to push the QoS down through the CMTS to the Cable Modem.

The DOCSIS network uses the concept of Service Flows to provide QoS. There are different types of DOCSIS Service Flows, such as Real Time Polling Service (RTPS) and Unsolicited Grant Service (UGS), which are tailored to carry different types of traffic. For upstream RTP traffic associated with VoIP, PacketCable recommends the use of Unsolicited Grant Service (UGS). In UGS, the CMTS issues a "grant" to the CM on a regular interval. In general, the size of the "grant" (i.e., the number of bytes to be transmitted by the CM when the grant is issued) is calculated based on the CODEC and 'p' time specified in the SDP offer/answer. The interval of the grants is calculated based on the 'p' time and the number of calls being multiplexed on a single UGS Flow. The CM uses these grants to transmit the RTP frames that match one of the classifiers associated with the UGS Flow. A classifier is comprised of a number of fields that uniquely identify an RTP session, such as source IP, destination IP, source port, destination port, and DSCP. Multiple classifiers can be associated with a UGS Flow. The CM compares the ingress traffic (Enterprise to SP network) against the classifiers assigned to UGS flows, and if a match is found the traffic is forwarded on the Service Flow associated with the classifier.

As discussed earlier, a UGS "grant" provides a fixed amount of bandwidth (i.e., a fixed number of bytes) to be transmitted by the CM in the upstream frames. If the packet (e.g., RTP packet) to be transmitted in the upstream direction is larger than the size of the UGS grant, the CM can not use the grant to transmit the packet. In this scenario, either the packet is discarded or, at implementers' option since this is unspecified in DOCSIS, sent using the best-effort service flow where it could also be discarded via rate limiting.

The QoS requirements for the ESG mostly revolve around ensuring that the upstream RTP packets match the interval, size, and DOCSIS classifier (From IP, To IP, From Port, To Port, DSCP) expected on the upstream UGS flow.

The ESG MUST override the DSCP field in the IP header of each upstream RTP packet to a configured value.

The ESG MUST override the DSCP field in the IP header of each upstream RTCP packet to a configured value.

The ESG MUST use a configured value for DSCP field in all upstream SIP messages received on the "Voice" port.

Since most PCMM Application Manager implementations assume an RTP header size of 12 bytes when calculating bandwidth for a voice call, the ESG MUST support the capability to modify the RTP header on upstream RTP packets to use the minimum-sized header, with the Extension 'X' bit cleared, and the CSRC Count 'CC' bits set to zero. The ESG MUST support a configuration data element that enables the operator to turn this feature on and off. Operators should enable this feature only after giving full consideration to the effects of removing RTP header extensions.

The ESG MUST support a configuration option where it filters [RFC 5761] requests to multiplex RTP and RTCP on the same port. If this filter this enabled, the ESG MUST modify the upstream SIP messages to not advertise the support for [RFC 5761].

The ESG MUST monitor the RTP upstream traffic to ensure that the CODEC and frame interval match what was negotiated in the SIP offer/answer. As a minimum requirement, the ESG MUST generate a vendor-proprietary log entry for each errant upstream RTP flow capturing the source IP address of the errant device. The ESG MAY choose to modify the DSCP field for upstream packets that are too large or arriving at an unexpected interval to the value '0'.

The ESG MUST detect RTP upstream traffic received on the "Voice" LAN port that is not associated with an active session and, as a configuration option, be able to filter these to prevent them from being transmitted on the PacketCable 2.0 access network.

The ESG MUST support a configuration option to allow RTP, RTCP, and SIP traffic received on the "Data" port to be transmitted on the PacketCable 2.0 access network.

An ESG with an embedded cable modem MUST shape traffic within the system so best-effort traffic on the "Data" port does not interfere with upstream RTP, RTCP, and SIP traffic on the "Voice" port.

6.3 Session Border Controller

The SBC provides SIP signal interworking, SIP-aware NAT/firewall, and IPv4/6 interworking functions for voice traffic between the Enterprise and Service Provider networks. The SBC can be loosely thought of as a single interworking rule set and a single firewall rule set that can be applied to one or more enterprise SIP entities. For example, if the SSP is providing hosted IP-Centrex service to multiple SIP phones that all support the same version of SIP, then these multiple phones could be served by an ESG containing a single SBC. If the SSP is providing SIP Trunking service to an enterprise containing two SIP-PBXs each supporting a different version of SIP, then the ESG would contain two SBCs each serving a single PBX. The number of SBCs instantiated in the ESG and how these SBCs map to the enterprise SIP entities depends on the deployment use-case, and how the operator chooses to configure the ESG to support that use-case.

Please refer to Annex A for a more formal description of the SBC object model.

6.3.1 Requirements

6.3.1.1 SIP Interworking

6.3.1.1.1 Configuring the Interworking Rules

The SBC MUST support a configuration option that enables the operator to select a predefined interworking rule-set that defines all the header manipulation and signal normalization rules for a specific SIP-PBX or SIP endpoint make and model. The SBC MUST also support a mechanism that enables the operator to add, remove, or modify individual header manipulation rules in the selected interworking rule-set.

If the enterprise SIP entity connected to pkt-esg-sig-2 complies with the SIP procedures defined for pkt-sig-esg-1, then the interworking rule-set is effectively "no rules". In this case where no interworking is required, and if SIP Digest authentication is disabled at the SBC as described in Section 9, then the SBC acts as a transparent pipe as far as SIP signaling is concerned. (Note, in this "transparent" mode, the ESG still performs other functions such as SIP-aware NAT, SIP-aware firewall, Telemetry.)

The SBC MAY support the following read-only parameters that are derived from the interworking rule set:

- SIP-Mode: Indicates whether the SBC is operating as a B2BUA, SIP Proxy, or does not appear as a SIP entity in the SIP signaling chain (e.g., where the enterprise SIP endpoints require no interworking rules).
- Service-Mode: Indicates whether the SBC is providing SIP Trunking interworking services where the enterprise SIP endpoint is a SIP-PBX, or hosted IP Centex interworking services where the enterprise SIP endpoint is a SIP phone.
- SIPConnectMode: When the SBC is providing SIP Trunking interworking services, this parameter indicates whether the SBC is operating in the "Registration Mode" or "Static Mode" defined in [SIPconnect1.1].

6.3.1.1.2 Mandatory SIP Procedures[7]

The SBC MUST support the Business SIP Services (BSS) Feature Specification [PKT-BSSF] on interface pkt-esg-sig-1 and pkt-esg-med-1.

The SBC MUST support SIPconnect1.1 Technical Recommendation [SIPconnect1.1] on interfaces pkt-esg-sig-1 and pkt-esg-med-1 when it is providing interworking services for SIP-PBXs over a SIP Trunking interface.

The specific set of procedures supported on pkt-esg-sig-1 and pkt-esg-med-1 for a given deployment is governed by the interworking rule-set.

This document does not place any requirements on the SIP/SDP signaling and media procedures supported on pkt-esg-sig-2 and pkt-esg-med-2.

6.3.1.1.3 Loop Detection

The SBC Interworking Function SHOULD detect SIP signaling loops. SIP signaling loops can occur due to call routing database errors, or call-forwarding loops. The SBC MUST preserve the History Info header information as specified in [RSTF] in signaling from the PacketCable 2.0 network. The SBC MUST preserve any History-Info header field received in signaling originated from the ESE. If no History-Info header field is received from the ESE, the SBC MUST insert a History-Info header field. The SBC MUST NOT reset the Max-Forwards header field or Max-Breadth header field [RFC 5393], or remove Via header fields.

[7] Updated by ESG-N-12.0682-8 on 10/28/16 by PO

6.3.1.2 SIP-Aware NAT/Firewall

6.3.1.2.1 SIP-Aware NAT

Each SBC MUST contain a SIP-aware NAT. The SBC NAT MUST maintain a table of IP address bindings that maps each LAN-side IP address and port to its equivalent WAN-side IP address and port. The SBC NAT MUST be capable of creating the table of address bindings dynamically, as it exchanges SIP messages between the Service Provider and Enterprise network. The SBC MUST also provide the capability to enable the Service Provider to provision the address bindings. When the address bindings are created dynamically, the SBC NAT SHOULD support [RFC 4787]. An address binding entry in the table contains two attributes: the ESEaddress data attribute containing the enterprise side IP address, and the UAaddress data attribute containing the public side IP address mapped to the enterprise IP address.

As it relays SIP messages and RTP/RTCP packets between the Enterprise and Service Provider networks, the SBC MUST map the IP addresses contained in IP headers, SIP headers, and SDP body that have both a LAN-side and WAN-side value such that the address translation is transparent to the Service Provider and Enterprise network. The SIP NAT MUST be capable of performing the NAT function for both IPv6 and IPv4 address types, including the case where IPv6 is used on one interface and IPv4 on the other.

The SBC SIP-aware NAT may be deployed in conjunction with an existing non-SIP-aware NAT/firewall located in an existing Enterprise network, or in a situation where there is no Enterprise NAT and the ESG provides a "one box" solution. When the SIP NAT is deployed with existing LAN infrastructure particular care must be taken to avoid placing a non-SIP-aware NAT between the ESG and the enterprise SIP endpoint, or between the ESG and the Service Provider network, since NAT traversal of external non-SIP aware NATs is not supported by the ESG. Also to take advantage of traffic shaping, QoS and other functions in the ESG and to enable SIP messages to be routed such that the SIP NAT will function properly, particular care must be taken when designing the non-SIP portions of the ESG for deployments into existing Enterprise network infrastructure.

6.3.1.2.2 SIP-Aware Firewall

The SBC MUST support a SIP-aware firewall (a.k.a. SIP firewall). The SBC SIP firewall MAY be configured to run in either promiscuous mode or firewall mode. In promiscuous mode all traffic is be passed without regard to source, destination, message type, rate or any other parameter. In firewall mode traffic MUST be filtered in accordance with the currently configured set of rules as indicated by the SIP-firewall-Rules data attribute.

The SBC MUST support a configuration option that enables the operator to select a predefined SIP firewall rule-set that defines all the SIP methods, headers, and response codes that are allowed. The SBC MUST also support a mechanism that enables the operator to add, remove, or modify SIP firewall rules in the selected SIP firewall rule-set.

The SBC MUST maintain an Access Control List (ACL) that contains an entry for each enterprise SIP endpoint in the Enterprise network, and each connected P-CSCF/IBCF in the PacketCable network. At a minimum the SBC MUST identify the IP addresses or domain name of these entities in their ACL entry. The SBC MUST provide a mechanism that enables the operator to view and modify the ACL.

The SBC SIP firewall MUST verify that SIP messages received from enterprise SIP entities are properly constructed and valid. If the SBC receives a SIP message that is malformed or not valid, then the SBC SIP firewall MUST take action in accordance with the configured SIP firewall rule-set.

On receiving a properly formed and valid SIP message from an enterprise SIP entity, the SBC SIP firewall checks the ACL to verify that the message originated from an entity that is allowed access. If the message originated from an entity that is not in the ACL, then the SBC SIP firewall MUST either silently drop the message, or send an error response, based on the configured firewall rule-set. If the SIP message originated from an authorized entity then the SBC SIP firewall MUST open the necessary ports to allow communication only for the duration of the call/dialog.

Therefore, to allow signaling messages through the firewall, the SIP firewall monitors the ports used by SIP signaling (typically port 5060 but provisionable). The SIP NAT/firewall extracts information and modifies fields in the messages to control the routing of signaling and media. The SIP NAT/firewall can dynamically open a pinhole in the firewall when a session establishment message is received, and close it upon completion of the call.

The SBC SIP firewall MUST be able to rate-limit SIP messages sent from the Enterprise SIP entities to the PacketCable 2.0 network. These include various keep-alive messages and REGISTER messages. This function helps protect the PacketCable 2.0 network from Denial of Service (DoS) attacks initiated from the enterprise network.

6.3.1.2.3 Firewall Logging Requirements

The SBC MUST log all events that are blocked by the SIP firewall, including invalid access attempts, malformed messages, rate limiting of SIP messages, etc. The SBC MUST be capable of logging the following SIP firewall events:

- All permitted inbound access requests for SIP/SDP and RTP from the public network.
- All permitted outbound access requests for SIP/SDP and RTP from private network clients that use the PacketCable 2.0 service.
- All dropped or denied access requests from private and public network that are meant to transverse the ESG that violate security policy.
- All dropped or denied access requests from private, service and public network to send traffic to the ESG itself that violate the security policy.
- All attempts to authenticate at an Administrative Interface on the ESG itself.
- All access requests from private, service and public network to send traffic to the ESG itself on the port or ports used for Remote Administration.
- Each startup; of the system itself or of the of the security policy enforcement component of the system.
- All manually entered changes to the system clock.

For each event logged, the SBC MUST capture the following log data elements:

- Date and Time - when the event occurred where the date recorded each event in the log consists of the four-digit year, the month and the date and the time recorded for each event in the log must consist of the hour, the minute and the second.
- Protocol as indicated in the IP header field.
- SIP Identity (if available)
- Source IP Address.
- Destination IP Address.
- Source Port (TCP and UDP [Transmission Control Protocol and User Datagram Protocol]).
- Destination Port (TCP and UDP).
- Message Type.
- Disposition of the Event.
- Statement of success or failure to authenticate at an Administrative Interface. Failed authentication attempts must include the reason for the failure.

6.3.1.2.4 RTCP ports & RTP ports opened per session

The Real-time Transport Protocol (RTP) is comprised of two components: a data transfer protocol (RTP), and an associated control protocol (RTCP). Historically, RTP and RTCP have been run on separate UDP ports. As specified in [RFC 4566], SDP identifies the RTP port number using the "port" sub-field of the Media Description "m=" line. The RTCP port number is either derived algorithmically from the RTP port, or it is explicitly identified in SDP using the "a=rtcp:" attribute.

When the RTCP port is derived from the RTP port, SDP uses the convention that RTP is assigned an even numbered port, and RTCP belonging to the same session uses the next higher odd port. A non-VoIP-aware NAT will typically destroy the ordering of ports in the translation process. The SBC SIP-aware NAT MUST preserve this port mapping convention.

6.3.1.2.5 Media Relay Requirements

The SBC MUST support a media relay function. The media relay function is not required to support media transcoding.

The SBC media relay MUST NOT modify the media packets (e.g., UDPTL, RTP) in either the upstream or downstream direction. The SBC media relay MUST NOT modify the RTCP packets in either the upstream or downstream direction. The purpose of these requirements is to ensure that the SBC media relay enables end-to-end negotiation of any type of media session that is supported by the enterprise and remote SIP endpoints. For example, in addition to basic G.711 voice, the SBC media relay is expected to support media sessions that carry other audio encodings such as compressed and wideband audio, other media types such as video, and voice-band data encoding schemes such as T.38 UDPTL, FAX, V.152, and DTMF-relay.

6.3.1.3 IPv4/6 Interworking

The SBC MUST support IPv4 to IPv6, and IPv6 to IPv4 interworking.

6.3.1.4 Administrative and Operational Status

6.3.1.4.1 Administrative State of SBC

The SBC MUST support management controls that enable the operator to remove it from service (say for routine maintenance or version upgrade). The SBC MUST support the following management commands:

- Remove from service immediately. On receiving this command, the SBC MUST initiate procedures to force-release all active calls and transition to an out-of-service state.
- Remove from service when number of active calls drops to zero. On receiving this command, the SBC MUST transition to an out-of-service state only after the number of active calls has dropped to zero. The SBC MUST block new call attempts while it is waiting for the active calls to release. If a "return to service" timer expires prior to transitioning to out-of-service state, then SBC cancels the command, returns to full service, and sends a management alert if provisioned to do so.

The SBC MUST support management controls that enable the operator to restore an out-of-service SBC to the in-service state.

The SBC MUST support a mechanism that enables these administrative commands to be scheduled at a predefined time.

Figure 11 shows the state transition diagram for the SBC administrative state.

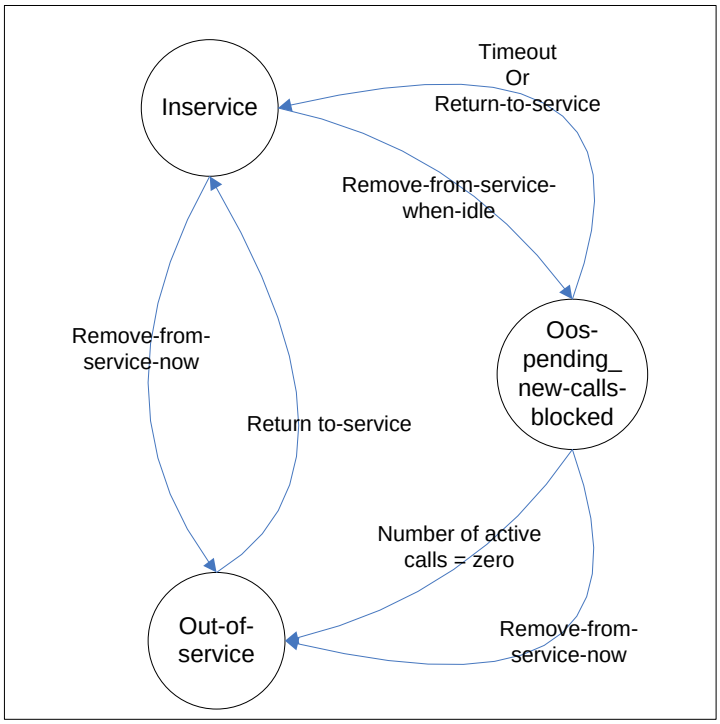

Figure 11 - SBC Administrative State Transition Diagram

The SBC MUST make the administrative status of the SBC available to the Service Provider.

6.3.1.4.2 Operational State of SBC[8]

The SBC MUST provide a mechanism that enables the operator to view its operational status; i.e., which indicates the SBC's ability to support Business Voice service to the Enterprise. This specification does not define the specific operational status values, but at a minimum they should indicate the following information:

- "operational" - the SBC is able to support Business Voice service
- "not operational" - some condition exists which prevents the SBC from providing Business Voice service.
- "survivable" - the SBC is partially operational. Specifically, the SBC has lost SIP connectivity with the Service Provider network, but is able to support intra-enterprise calls to the enterprise SIP entities over its voice LAN interface as described in section 6.3.1.5. For the case where multiple SIP signaling connection points are configured for this SBC, the loss of SIP signaling connectivity means that the SBC has lost connectivity with *all* connection points.

An SBC in the "operational" state MUST transition to the "survivable" state on detection of loss of SIP connectivity with the Service Provider network.

Note: Loss of SIP connectivity with the Service Provider network can be caused by many different fault conditions. The fault condition could be local to the ESG and therefore explicitly detected by the ESG (e.g., loss of the physical WAN interface). Or the fault condition could be upstream from and not directly detectable by the ESG (e.g., loss of layer-2 or layer-3 IP connectivity between the CMTS and the P-CSCF, or loss of the SIP signaling interface on the P-CSCF).

[8] Updated by ESG-N-12.0682-8 on 10/28/16 by PO

For fault conditions that are explicitly detected by the ESG, such as loss the physical WAN interface, the SBC MUST transition immediately to the "survivable" state on onset of the fault condition. How this requirement is applied varies between the S-ESG and the E-ESG.

1) For an S-ESG, the SBC MUST transition to the “survivable” state immediately on detection of loss of layer-1 on the WAN Ethernet port.

2) For an E-ESG, loss of the physical WAN interface is equivalent to loss of the DOCSIS access network. The availability status of DOCSIS is tracked by the eCM eSAFE entity within an E-ESG, and therefore the eCM MUST inform the eESG of the availability status of the DOCSIS access network (how this information is signaled is not specified). An SBC within an eESG MUST transition to the “survivable” state immediately on being informed by the eCM that the DOCSIS access network is not available.

For fault conditions that cannot be explicitly detected by the ESG, the SBC MUST provide a mechanism to implicitly detect the fault (e.g., an OPTIONS ping). The mechanism to implicitly detect loss of SIP connectivity SHOULD be sufficiently robust that it does not interpret a transient network failure as loss-of-connectivity. For example, a transient overload condition that temporarily delays delivery of SIP messages or that drops a single SIP message should not be interpreted as a loss of connectivity. The mechanism to implicitly detect loss of SIP signaling MUST be capable of detecting hard connectivity failures (e.g., catastrophic failure of the P-CSCF) within 1 minute of the onset of the failure. The SBC MUST provide a configuration control that enables the operator to adjust the responsiveness of this mechanism to detect hard failures (e.g., a configuration parameter to control the OPTIONS ping interval).

When the fault condition that drove the SBC into the "survivability" state clears and SIP signaling connectivity is restored, and if there is no other condition preventing the SBC from providing service, then the SBC MUST transition to the "operational" state.

There can be multiple conditions that could cause an SBC to become "not operational". For example, the SBC could be missing some critical configuration data, or a physical failure such as loss of the voice LAN port could be blocking its ability to provide service to the enterprise SIP entities.

When it is in a non-operational state, the SBC MUST provide a mechanism to convey to the operator the specific condition that is preventing it from becoming operational.

6.3.1.4.3 Administrative State of Enterprise SIP Entity

The SBC MUST support a configurable state parameter that represents the administrative state of an enterprise SIP entity. An enterprise SIP entity can be administratively removed from service for two reasons; because the entity itself has been administratively removed from service via the management interface, or because its super-ordinate SBC has been administratively removed from service. The SBC MUST make the administrative status of the SIP entity available to the Service Provider. The SBC MUST convey sufficient state information to indicate why the SIP entity is administratively out-of-service (i.e., whether the entity itself has been removed from service or the super-ordinate SBC has been removed from service).

When a SIP entity that represents a hosted SIP terminal is removed from service, the SBC SHOULD attempt to release all active calls and de-register the entity from the SP network as specified in [PKT 24.229].

When a SIP entity representing a SIP-PBX is removed from service, the SBC SHOULD release all active calls toward the SP network. If the SBC is configured to operate in the SIPconnect1.1 "registration mode" on pkt-esg-sig-1, then it MUST also de-register the entity from the SP network as specified in [SIPconnect1.1].

The SBC MUST ignore any SIP message received on pkt-esg-sig-1 for a SIP entity that has been administratively removed from service.

6.3.1.4.4 Operational State of Enterprise SIP Entity

The SBC MUST provide a mechanism to detect if the connected Enterprise SIP entity (SIP-PBX or SIP endpoint) is operational. If the connected Enterprise SIP entity supports registration, then the SBC MUST base the operational status on the registration state of the connected device, where "registered" means the device is operational, and "not registered" means it is not operational. If the connected Enterprise SIP entity does not support registration, then the mechanism to determine its operational status is not specified (e.g., the SBC could use OPTIONS ping).

In addition to the "operational" vs. "not operational", the SBC MAY support additional sub-states, e.g., "device is registered but not responding to incoming requests."

The SBC MUST make the operational status of the Enterprise SIP entity available to the Service Provider. When the ESE is in a non-operational state, the SBC MUST provide a mechanism to convey to the operator the specific condition that is preventing the ESE from becoming operational.

6.3.1.5 Survivability[9]

Section 6.3.1.4.2 describes how the ESG tracks the SIP signaling connectivity status of the WAN interface. This section (Section 6.3.1.5) extends those requirements to describe how the ESG supports intra-enterprise calls when it loses SIP signaling connectivity with the Service Provider network.

This section uses the term “survivability” to describe the general capability where an ESG that is isolated from the Service Provider network acts on its own accord to provide basic call service to the enterprise phones. The primary survivability use case applies to hosted IP-Centrex service, where the ESG supports abbreviated-dialed calls (e.g., 4-digit extension calls) among the hosted SIP phones within a single enterprise. However, the procedures described here are sufficiently general to support survivability for any service offering, such as an intra-enterprise calls to both E.164 and abbreviated extension numbers for an enterprise that has multiple SIP-PBXs and multiple hosted SIP phones, all served by the same ESG.

The remaining subsections within this section describe the procedures that must be supported by the SBC as it transitions between the "normal" (aka "operational") and "survivable" states.

6.3.1.5.1 SBC Behavior While in Normal Mode

While in "normal" mode, the SBC MUST maintain LAN location binding information for the registered enterprise SIP entities; i.e., the binding between the Public User Identities of an enterprise entity and the LAN contact address of the entity that is maintained by the ESE(lan) object defined in the ESG object model in Annex A (see A.1.1.5 for more details).

6.3.1.5.2 SBC Behavior in Transition Into Survivability

When the SBC transitions into the "survivable" mode, it MUST maintain all stable 2-way intra-enterprise calls until they are released though normal user activity (e.g., a user hangs up). For example, on entry into survivable mode, the SBC could take on the role of a B2BUA to associate the dialog pair of each stable 2-way intra-enterprise call, routing mid-dialog requests between the endpoints, and releasing the dialogs when the call ends.

Note: A call is considered to be a stable 2-way enterprise call if the call is between two SIP endpoints in the enterprise, and the media for the call is within the enterprise LAN network (i.e., the RTP source and destination IP addresses are both within the LAN address space). By this definition, a call in the "ringing" state would be considered stable if the endpoints have exchanged media addresses (e.g., in the SDP exchanged in an INVITE request and 18x response).

[9] Section added by ESG-N-12.0682-8 on 10/28/16 by PO

For dialogs that it does not maintain (e.g., dialogs associated with calls between the enterprise and Service Provider network) the SBC MUST ensure that the dialog is released gracefully. For example, the SBC can take on the role of the remote SIP User Agent for these dialogs and release them using the appropriate SIP procedure (e.g., by sending a BYE request to the enterprise endpoint to release a dialog associated with an answered 2-way call). In this context, an active SIP registration is not considered a dialog.

Note: An implementer could choose to release non-intra-enterprise calls immediately. Or, an implementer could decide to defer call release until the user hangs up to account for the failure case where SBC loses SIP signaling connectivity but still has RTP connectivity with the Service Provider network.

This specification does not place any call-handling requirements on the SBC on entry into survivable mode beyond the requirement to maintain stable 2-way intra-enterprise calls. As an implementation option, the SBC could maintain additional call types such as stable intra-enterprise 3-way calls (say, for the case where the media mixing function is performed by one of the enterprise endpoints in the 3-way call), or dialogs associated with non- call activity (e.g., SUBSCRIBE dialogs associated with shared call appearances).

6.3.1.5.3 Behavior While in Survivable Mode

While in the "survivable" mode, the SBC MUST maintain the LAN location bindings between the enterprise user identities and users' endpoint LAN contact addresses. These bindings include the bindings that were learned before the SBC entered the "survivable" mode, minus any bindings that are removed due to enterprise SIP entities de-registering.

On receiving a new-dialog INVITE from a registered enterprise SIP entity served by the SBC, the SBC MUST recognize whether or not the INVITE Request-URI identifies another enterprise SIP entity served by the ESG. If the Request-URI identifies a registered SIP entity served by the ESG, then the SBC MUST forward the INVITE to that SIP entity using the location-binding information contained in the target ESE(lan) object, and follow the normal procedures to establish a 2-way call.

If the Request-URI identifies a user that is not served by the ESG, or that is served by the ESG but is not registered, then the SBC MUST reject the INVITE by providing a local announcement and/or sending a 4xx response.

6.3.1.5.4 Transition Out-of Survivability

When the SBC transitions out of "survivable" mode, it MUST maintain the currently established intra-enterprise calls as if it was still in the "survivable" mode (i.e., the SBC must handle in-dialog messages for these dialogs locally, and not forward them to the SP network). Also, the SBC MUST relay the next REGISTER request it receives from each enterprise SIP endpoint to the Service Provider network.

Note: Normally, the operator configures the Service Provider network and the ESG with a short registration refresh interval on the LAN side (say, 30 seconds), and a long refresh interval on the WAN side (say 30 minutes). On exit from survivability, the ESG should always send the next REGISTER request it receives from the enterprise SIP endpoint independent of where it is on the WAN registration refresh cycle. This will ensure that registration bindings for the enterprise SIP endpoints are resynchronized with the Service Provider network in a timely manner.

6.4 Telemetry

This section provides common requirements for the voice statistics that are applicable to the ESG. In addition, it specifies the common ESG requirements for event logging and reporting that are related to voice statistics and SIP messaging.

6.4.1 Requirements

6.4.1.1 VoIP Metrics

As the ESG acts as a media relay element between the SIP-PBX endpoint and the remote SIP endpoint, all RTP and RTCP traffic will traverse the ESG. The ESG is, therefore, in a position to gather the following statistics related to network performance:

1. Packet Interarrival Jitter: The ESG MUST calculate the Packet Interarrival Jitter as specified in [RFC 3550]. The jitter value is smoothed as in [RFC 3550] and relates to a smoothed jitter estimate since the beginning of the RTP session. The ESG MUST calculate Packet Interarrival Jitter for both upstream and downstream directions. The upstream Packet Interarrival Jitter is measured using the RTP packets received from a SIP-PBX endpoint and represents the jitter in the packet stream from the SIP-PBX endpoint to the ESG. The downstream Packet Interarrival Jitter is measured using the RTP packets received from a remote SIP endpoint and represents the jitter in the packet stream from the remote endpoint to the ESG.
2. Packet Loss. The ESG MUST calculate the following two packet loss estimates:
    a. The fraction of RTP packets lost during a configurable time interval Tpl, where Tpl has a default value of 5 seconds. This corresponds to the 'fraction lost' calculation in [RFC 3550], except the period of calculation is a configurable period rather than since the last RTCP report as in [RFC 3550]. The timer corresponding to Tpl is started at the beginning of RTP transmission and the count of the fraction of RTP packets lost is reset every Tpl seconds.
    b. The cumulative number of RTP packets lost since the beginning of reception as per [RFC 3550]. As in [RFC 3550], the number of packets lost does not include late or duplicate packets.
3. The ESG MUST calculate the above two RTP packets loss estimates for both upstream and downstream directions. The upstream Packet Loss is measured using the RTP packets received from a SIP-PBX endpoint and represents the loss in the packet stream from the SIP-PBX endpoint to the ESG. The downstream Packet Loss is measured using the RTP packets received from a remote SIP endpoint and represents the loss in the packet stream from the remote endpoint to the ESG.
4. Round-Trip Propagation Delay. If remote RTCP sender reports are available, the ESG MUST measure the following two Round-Trip Delays, with both delays relying on the sending of RTCP reports by the SIP endpoints:
    a. The delay related to the leg from the ESG to the SIP-PBX endpoint and back to the ESG.
    b. The delay related to the leg from the ESG to the remote SIP endpoint and back to the ESG.
5. Both these delays can be calculated as in [RFC 3550]. Both Round-Trip Delay values MUST be calculated each time an RTCP SR report is received. It should be noted that the delay calculation is dependent on accurate reporting of parameters such as the DLSR by the SIP endpoints. As the RTP/RTCP implementation in the SIP endpoints (particularly the SIP-PBX endpoint) cannot be guaranteed to be accurate, the ESG SHOULD implement some error checking to ensure that patently erroneous values of delay are not reported.
6. Burst/Gap Parameters. The ESG MAY measure the following burst/gap parameters: burst loss density, burst duration, gap loss density and gap duration. These parameters are calculated as in [RFC 3611] for both upstream and downstream directions. The upstream burst/gap parameters are measured using the RTP packets received from the SIP-PBX by the ESG. The downstream burst/gap parameters are measured using the RTP packets received from a remote SIP endpoint.

Although standard RTCP packets as well as possibly RTCP-XR packets pass through the ESG, it is assumed that the ESG may inspect certain parameters from these packets as necessary, but it MUST NOT modify the contents of the RTCP or RTCP-XR packets.

If RTCP sender reports are available from the remote endpoints, the ESG MAY report the 'cumulative number of packets lost' and 'interarrival jitter' specified in [RFC 3550] for both the upstream and downstream directions.

If RTCP-XR reports are available from the remote endpoints, the ESG MAY report the parameters associated with the RTCP-XR VoIP metrics block specified in [RFC 3611] for both the upstream and downstream directions.

6.4.1.2 Call Statistics

At the end of each call, the ESG MUST, at a minimum, measure, collect and store the following statistics for that call:

- Call Start Time
- Call End Time
- SIP Call-ID
- Direction (Inbound to SIP-PBX or Outbound from SIP-PBX)
- Originating and Terminating SIP URIs
- Codec Type.

6.4.1.3 SIP PUBLISH Mechanism

The SIP PUBLISH mechanism for the reporting of RTCP-XR VoIP metrics from the ESG to a performance management function located in a back-office server is specified in [RFC 6035]. The back-office server function that receives the VoIP metrics reports is referred to as the 'Telemetry Collector' device. This is typically an element manager or network manager that is responsible for VoIP session/media performance management. During registration, the ESG MUST indicate support of the vq-rtcpxr package defined in [RFC 6035]. It is informed of the contact address of the Telemetry Collector as part of the provisioning process. The ESG MUST populate the request URI in the PUBLISH request with the Telemetry Collector address. The ESG MUST send separate PUBLISH messages for the upstream leg of the call towards the remote SIP endpoint and for the downstream leg of the call towards the SIP-PBX endpoint. The ESG MUST populate the RemoteID parameter within the PUBLISH message body with the SIP URI of the remote SIP endpoint for both the upstream and downstream legs of the call to allow the operator to determine to which direction each PUBLISH message relates.

As the ESG is acting as a RTP/RTCP media relay rather than an RTP/RTCP endpoint, only certain RTCP-XR VoIP metrics can be calculated and sent in the 'LocalMetrics' block of the PUBLISH messages to the Telemetry Collector. The following parameters MUST be included in the 'LocalMetrics' block:

- NetworkPacketLossRate (NLR)
- RoundTripDelay (RTD)
- InterarrivalJitter (IAJ).

The following parameters MAY be included in the 'LocalMetrics' block:

- BurstLossDensity (BLD)
- BurstDuration (BD)
- GapLossDensity (GLD)
- GapDuration (GD)
- MinimumGapThreshold (GMIN).

The ESG MUST set the parameter GMIN to the value 16. The ESG MUST report the parameters calculated above in separate PUBLISH messages for both the upstream and downstream legs of the call.

The ESG MAY also populate the 'RemoteMetrics' block of the PUBLISH message from RTCP/RTCP-XR reports that are sent by the remote SIP endpoint or the SIP-PBX endpoint if these reports are available.

6.4.1.4 Reporting Requirements

The ESG MUST support the following mechanisms to make the VoIP metrics and call statistics information available to the operator:

- Local log accessible from a Web GUI
- Sending VoIP metrics and call statistics to an external syslog server.

The ESG MUST make these mechanisms available for reporting statistics for both (a) the upstream leg of a call between the ESG and the remote SIP endpoint and (b) the downstream leg of the call between the ESG and the SIP-PBX endpoint.

The ESG MUST store VoIP metrics and call statistics in the local call log for at least the most recent Ncr calls that have traversed the ESG, where Ncr is a configurable attribute indicating the Number of Calls to be Recorded in the log. Ncr has a default value of 10.

At the end of each call, the ESG MUST send VoIP metrics and call statistics related to that call to an external Syslog server, when configured to do so. The message will be sent on Facility value 16 (local use 0) and severity 6 (informational) to give a Priority value of 134.

In addition, the ESG SHOULD provide SNMP configurable alarm thresholds on network performance and audio quality metrics (if available via direct measurement or RTCP reports inspection) that will generate an SNMP trap.

There are three types of metrics reports that use the SIP PUBLISH mechanism: session reports, interval reports and alert reports. The ESG MUST support session reports and, when configured to do so, report metrics to the Telemetry Collector at the end of each session or when a media change occurs. In addition, the ESG MAY support interval reports and alert reports when the ESG is placed into an 'RTCP-XR VoIP metrics debug mode'. Once the ESG is placed into this mode, mid-call interval reports will be sent to the Telemetry Collector on a regular basis. Alert reports will also be sent to the Telemetry Collector once the ESG is in debug mode if pre-configured quality thresholds are breached for the RTCP-XR VoIP metrics being reported to the Telemetry Collector. These reports and alerts are generated for all active sessions or for a particular active session, according to provisioning. As noted in [RFC 6035], care should be employed to avoid overload when placing the ESG into the RTCP-XR VoIP metrics debug mode as it is possible that large numbers of PUBLISH messages will be sent by the ESG to the Telemetry Collector. The debug mode should, therefore, be employed as a temporary means of troubleshooting rather than a normal mode of operation.

6.4.1.5 SIP/RTP Tracing

The ESG MUST capture and store SIP signaling traces per call for at least the most recent Ncr calls that have traversed the ESG. Ncr is configurable and has a default value of 10. In each of these traces, the SBC MUST include:

- Unique Trace ID
- Call Start Timestamp
- SIP Call-ID
- Originating and Terminating SIP URIs
- All SIP messages up to the current time for the corresponding call.

For each trace, the ESG MUST show both the signaling between the ESG and the SIP-PBX and the ESG and the PacketCable 2.0 network on the same trace with all messages presented in strict chronological order. SIP signaling traces MAY be presented as a ladder diagram accessible from the Web GUI of the ESG. In addition, the ESG MUST support the upload of the SIP traces to an external FTP server.

Triggered by the operator, the ESG MUST be able to capture and store at least the next Tcr duration of all RTP media streams that traverse the ESG, with Tcr being configurable and having a default value of 0 second. This functionality can be used for diagnostics, audio playback and troubleshooting purposes. The ESG MUST support the upload of the RTP media streams to an external FTP server. In addition, the ESG SHOULD provide the ability to filter the traffic capture based on signaling packets, RTP packets for a particular SIP URI or IP address.

For the convenient viewing of the SIP signaling and RTP media traces by the operator, the captured packets MAY be stored in the data format defined for a popular capturing tool such as WireShark.

6.5 SIP Endpoint Test Agent (SETA)

6.5.1 Overview

As discussed in the previous sections, the Enterprise SIP Gateway (ESG) operates as a demarcation point between the Service Provider and Enterprise networks. One of the main functions of the ESG is fault detection and isolation. SETA plays a role in this function by providing a management interface that enables the operator to initiate test calls from the ESG to a pre-programmed termination point in the Service Provider network, either on demand or at a programmed periodic interval. The SETA can also accept test calls, including RTP loopback calls (RTP packet reflector only). Data collected from the SETA test calls can be used by the Service Providers to take pre-emptive action in resolving problems before they become visible to the customer.

SIP Endpoint Test Agent (SETA) is a logical function within the ESG which can perform the fault isolation and probing functions via management commands, and without any interaction with the enterprise users. SETA performs its functions independent of the service or operational state of the enterprise SIP endpoints.

Section 6.5.2 provides the technical details that the ESG must support to implement a SETA function compliant to this specification. Figure 12 and Figure 13 identify the administrative demarcation point between the Service Provider and Enterprise network. These diagrams are for the illustration purposes only, and are not meant to place any restrictions on how operators manage and maintain their business service offerings.

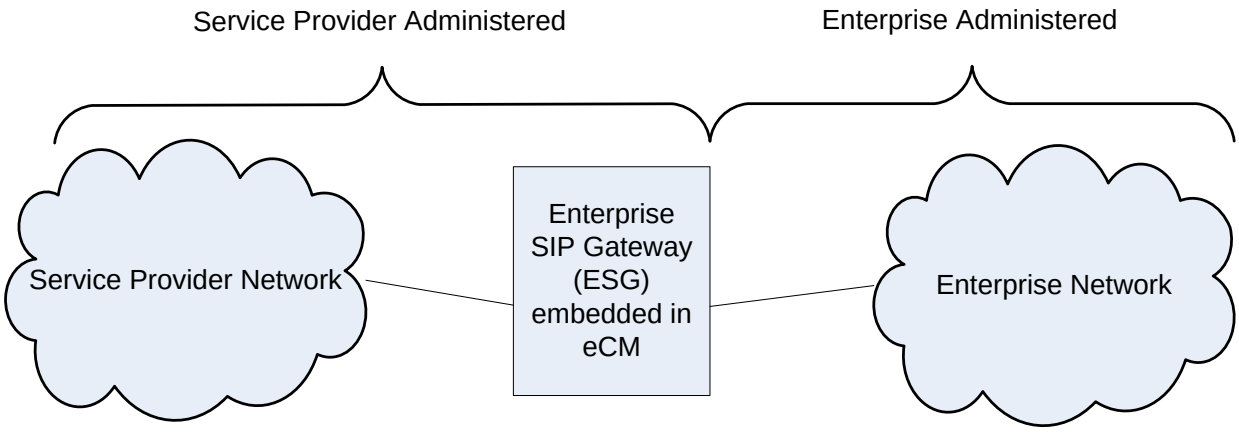

Figure 12 - Administrative Demarcation Point for Embedded ESG

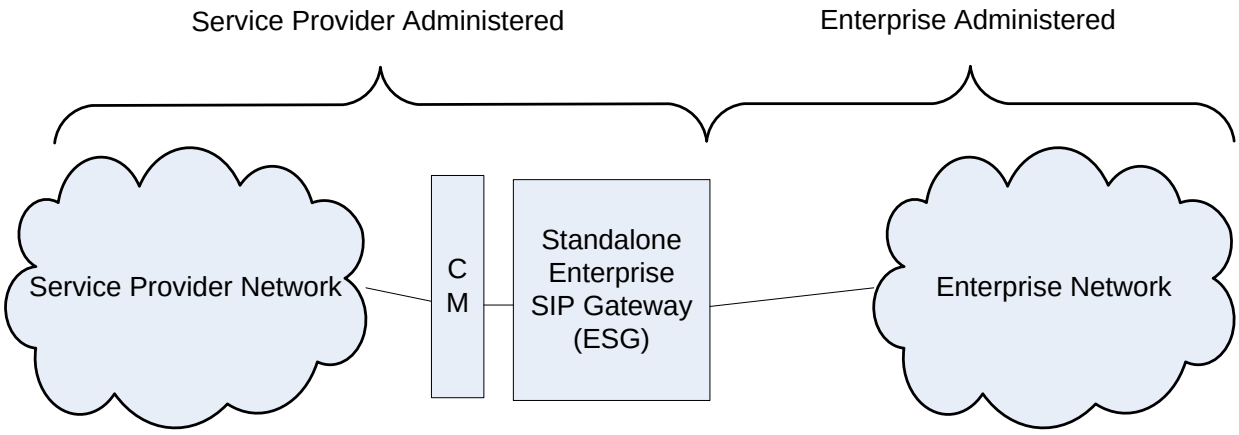

Figure 13 - Administrative Demarcation Point for Stand-Alone ESG

6.5.2 Technical Requirements

6.5.2.1 General Requirements[10]

The ESG MUST support a SETA function. The ESG MUST provide management controls to enable and disable the SETA functionality. Since the ESGs are expected to work in both "Registration" and "Static" mode as defined in [SIPconnect1.1], SETA MUST also support both modes. Specifically, when operating in Registration mode, SETA appears as a PacketCable 2.0 compliant registering SIP endpoint connected via the P-CSCF to the Service Provider network. When operating in the static mode, SETA appears as an endpoint in a peer network connected via the IBCF to the Service Provider network. SETA MUST support a configuration option which allows the operator to change the mode of operation as required.

SETA is identified by its assigned IP Multimedia Public Identity (IMPU). The ESG MUST ensure that calls targeted to this IMPU are routed to SETA. The ESG MUST ensure that SIP signaling and RTP messages for all other calls are not routed to SETA.

When it appears as a SIP endpoint hosted by the PacketCable 2.0 network, SETA MUST support the basic call originating and terminating procedures defined in [PKT 24.229].

When it appears as a SIP-PBX, SETA MUST support the basic call originating and terminating procedures defined in [SIPconnect1.1].

As discussed in the overview section, SETA is controlled solely by the management interface (e.g., SNMP) and does not require input from or cooperation of the enterprise user. Considering the important functions that SETA performs and potential impacts to the Service Provider network if not used appropriately, the ESG MUST NOT expose the SETA management interface to the enterprise.

6.5.2.2 Requirements when the ESG is working in the Registration Mode

When configured in the Registration mode, SETA MUST register to the operators' network and follow the UE (User Equipment) requirements defined in [PKT 24.229]. SETA MUST register to the network independent of the state of the enterprise SIP endpoints. After successful registration, SETA MUST support the UE call origination and termination procedures defined in [PKT 24.229]. The ESG MUST reject (e.g., ICMP port unreachable), or ignore incoming calls targeted to SETA, if the SETA is not registered.

[10] Updated by ESG-N-12.0682-8 on 10/28/16 by PO

Based on local configuration, SETA can register to the Service Provide Network using:

- the same P-CSCF and TLS (Transport Layer Security) connection used for the enterprise SIP endpoints, or
- the same P-CSCF, but via a different TLS connection than that used for the enterprise SIP endpoints, or
- a different P-CSCF and TLS connection than that used for the enterprise SIP endpoints.

For example, by configuring SETA with the same IMPI (IP Multimedia Private Identity) and the same P-CSCF as the enterprise SIP endpoint, the operator can cause SETA to use the same P-CSCF and TLS as the enterprise SIP endpoint.

Note: A Service Provider can use this configuration to not only verify the connectivity and voice quality between the Service Provider network and the ESG, but also to check the signaling and media path for an individual enterprise SIP end point with same IMPI as SETA.

6.5.2.3 Requirements when the ESG is working in the Static Mode

When configured in Static mode, SETA MUST not register to the operators' network. Instead, the SETA routing information is configured in the Service Provider network.

The ESG MUST reject (e.g., ICMP port unreachable) or ignore the incoming calls, targeted to SETA, if the SETA is not enabled.

6.5.2.4 Test Call Termination Requirements

The Service Provider network may initiate a test call to the ESG. This allows the network operator to test the voice path and to collect a number of voice/packet performance statistics.

The ESG MUST receive and terminate calls targeted to SETA from the Service Provider. The ESG MUST not send any protocol (e.g., SIP, RTP and RTCP) messages associated with these calls towards the enterprise network. Additionally, the ESG MUST reject or silently discard any incoming calls targeted for SETA from the customer network. As part of the call termination, SETA, MUST support the following two modes of operations:

- **RTP Packet Loopback:** SETA MUST support the 'rtp-pkt-loopback' mode as an answering entity as defined in [PKT-RST-E-DVA]. [PKT-RST-E-DVA] defines two sub-modes for 'rtp-pkt-loopback': "encapsulated RTP", and "payload loopback". SETA MUST support "payload loopback". SETA MAY support the "encapsulated RTP". SETA is not required to support the 'rtp-media-loopback' and 'rtp-start-loopback' mode as an answering entity as defined in [PKT-RST-E-DVA]. The offering entity (e.g., the test tool) in the Service Provider network is expected to follow the encoding requirements detailed in the [PKT-RST-E-DVA] specification.
- **Auto Answer:** If the incoming call targeted to SETA is a normal call (e.g., not requesting RTP loopback), then SETA MUST auto answer the call by responding with 200 OK to the INVITE. Additionally, SETA MUST follow the SDP Offer and Answer requirements in [RFC 3264] and [PKT 24.229]. In addition to auto answering the call, SETA MUST play the content of a stored audio file on the established media session to the remote endpoint. The SETA MUST be able to play the content of the file based on the locally configured CODEC and ptime, and the CODEC information received in the SDP from the call originator. SETA MUST play the content of the file over and over until either: 1) the call is terminated by the call originator, or 2) the time Call-Length expires at the ESG.

For efficient conduct of quality test, an audio file with a length less than 5 seconds is not recommended. Additionally, it is recommended that the content of the audio files should be speech-like where the speech-to-silence ratio is at least 1 or higher. Finally, the Call-Length should not be configured with a value less than 50 seconds to ensure the test call lasts long enough to enable accurate measurement of voice quality metrics. If the call duration is less than 50 seconds, VoIP metrics are not likely to be sufficiently accurate (particularly delay measurements which rely on receiving RTCP reports from the far end). This specification recommends a default Call-Length value of 300 seconds.

For both the RTP Packet Loopback and Auto answer scenario, SETA MUST terminate the call sending a 200 OK response on receipt of a BYE request from the call originator. In the absence of a BYE request from the call originator, SETA MUST terminate the call by sending a BYE request to the call originator when the Call-Length timer expires.

SETA is required to support only one active session at a time. If SETA receives an incoming call when it is establishing a session or already in an active call, it MUST reject the incoming call with 486 busy and continue with the other session normally.

6.5.2.5 Test Call Origination Requirements

The Service Provider may request the ESG to initiate a test call towards the Service Provider network. Test call origination at the ESG provides an additional tool that enables the Service Provider to test the voice path, and collect voice/packet performance statistics. The Service Provider can use this functionality to test the connection and collect statistics between the ESG and the core network, and also between two endpoints at the edge of their network (e.g., between SETA and an enterprise SIP endpoint).

In order to originate a call from SETA, the Service Provider instructs SETA (using SNMP or other applicable management protocol) to make call(s) to a specific remote endpoint. At a minimum the Service Provider provides the address of the remote endpoint (e.g., URI), length of the call, and the audio file to be played towards the remote end point. If the Service Provider wants to use this feature to make calls on reoccurring basis, the Service Provider also provides the reoccurring schedule (e.g., number of calls per 24 hrs.). The reoccurring schedule can be configured using the "call schedule" parameter. Finally, the Service Provider activates the feature which triggers SETA to originate a single call or multiple reoccurring calls.

When instructed via the provisioning interface to initiate a test call, SETA MUST establish an outgoing call per the procedures specified above in Section 6.5.2.1.Once the test call is established, SETA MUST play the contents of the audio file over and over until either:

1) the call is terminated by the remote endpoint, or
2) the Call-Length timer expires.

SETA MUST terminate the call with 200 OK on receipt of a BYE request from the remote endpoint. In the absence of a BYE from the remote endpoint, SETA MUST terminate the call by sending a BYE request when the Call-Length timer expires.

The audio file and Call-Length parameters for the originating test call should follow the recommendations provided for the audio file and Call-Length parameters in the Test Call Termination Requirements Section 6.5.2.4.

To prevent unintended calls, the ESG SHOULD by default disable the SETA call originating function. The ESG MUST provide provisioning controls that enable the operator to enable the SETA call originating function using configuration file or management commands (e.g., SNMP set).

6.5.2.6 Call Statistics Collection and Reports

SETA MUST collect statistics for both Signaling and Media. The details of the statistics to be collected and how they should be reported are provided below.

6.5.2.6.1 SIP statistics Collection and Reports

The SETA MUST capture and store the following statistics for at least the most recent Ncr calls that were either originated from or terminated to SETA. (Note, this Ncr call count is specific to SETA, and is separate from and independent of the Telemetry Ncr call count described in 6.4.1.5.) The SETA Ncr value is configurable, with a default value of 10.

The SETA MUST collect the SIP call statistics as described the Telemetry Section (6.4) of this document.

For the case where the call originates from SETA, SETA MUST collect the following additional statistics:

- Session Request Delay (SRD) as per [RFC 6076]
- Session Disconnect Delay (SDD) as per [RFC 6076]
- SIP Response code received for the INVITE
- Call completion status

For the case where the call terminates at SETA, SETA MUST collect the following additional statistics:

- Session Disconnect Delay (SDD) as per [RFC 6076]
- SIP Response code sent for the received INVITE
- Call Completion status

Once call statistics collection is complete, the Telemetry function is responsible for reporting the data. The Telemetry function MUST support the following mechanism to make the SIP statistics information available to the operator:

- Local log accessible from a Web GUI
- Syslog

6.5.2.6.2 RTCP and RTCP XR VoIP Metrics Requirements

The SETA, as a call originating and terminating entity, MUST support the RTCP requirements defined in the section 7.2 of [PKT-CODEC-MEDIA] with following qualifications. The SETA MAY support the ability to disable the RTCP on a per session basis. The SETA MUST support the Network Performance Measurements as described in the section 7.8.1 of [PKT-CODEC-MEDIA]. The SETA MAY support the Audio Quality Measurements as described in section 7.8.1 of [PKT-CODEC-MEDIA].

Once VoIP metrics collection is complete, the Telemetry function is responsible for reporting the VoIP metrics data. The Telemetry function MUST support the following mechanisms to make the RTCP and RTCP XR VoIP metrics information available to the operator:

- Local log accessible from a Web GUI
- SIP PUBLISH
- Syslog

When reporting VoIP Metrics for calls to and from SETA using SIP PUBLISH, the ESG MUST conform to the UE requirements in section 7.8.1.1 of [PKT-CODEC-MEDIA] and SIP PUBLISH requirements in the telemetry section 6.4 of this specification. Since the SETA line is the originator and terminator of the SETA test calls, the concept of upstream and downstream leg does not apply to the SETA calls.

6.6 Data NAT/Firewall[11]

The ESG SHOULD support a Data NAT/Firewall function. If the ESG supports the Data NAT/Firewall function, it MUST comply with the Data NAT/Firewall requirements specified in this document.

[11] New sections added by ESG-N-12.0691-8 on 11/4/16 by PO

The ESG Data NAT MUST support the Traditional NAT functionality as defined in Service Provider network. In addition, the ESG Data NAT MUST support the Twice NAT functionality defined in [RFC 2663] to transparently route TCP and UDP packets between an Enterprise SIP endpoint in the Enterprise network and a server in the Service Provider network. In addition, the ESG Data NAT MUST support the Twice NAT functionality defined in [RFC 2663], where the NAT interworks the destination IP address:port of TCP or UDP packets received from the ESE based on locally configured LAN→WAN IP address:port mapping information. The ESG MUST provide a configurable option to enable or disable this Twice NAT functionality. The ESG Data NAT SHOULD support the traditional NAT requirements specified in [RFC 4787].

The ESG MUST support a general-purpose Data Firewall bridging the WAN and LAN networks, along with configuration controls that enable the operator to define a rule set specifying which packets are allowed between the WAN and LAN networks.

7 DEVICE REQUIREMENTS[12]

This specification introduces two types of ESG devices. One is an embedded device that integrates an ESG function with a DOCSIS Cable Modem in the same device. The other is a stand-alone device that includes the ESG functionality but not a Cable Modem. The stand-alone ESG is connected to the PacketCable network via an external router or gateway (e.g., Cable Modem).

Both embedded and stand-alone ESGs MUST support the requirements in Section 6 of this document.

All flavors of the ESG MUST support at least one IP address when configured in the single-stack mode, and at least two IP addresses (at least one IP address per IP version) when configured in the dual-stack mode, for the interface facing the Service Provider network. This specification refers to this interface as the "WAN-side interface", and to the IP address(es) as the "WAN IP address(es)". The ESG MUST follow the procedures in Section 8 to obtain its WAN IP address(es). The ESG uses the WAN IP address(es) to communicate with the Service Provider network.

All flavors of the ESG MUST support at least one IP address on the interface facing the Enterprise network administered by the customer. This specification refers to this interface as the "LAN-side interface", and to the IP address as the "LAN IP address". The specific procedures to obtain the LAN side IP address are out of scope of this specification. The ESG may allow the customer to configure a static LAN IP address, or it may obtain the LAN IP address from an enterprise DHCP server. The ESG uses the LAN IP address to communicate with the Enterprise SIP endpoints in the Enterprise network.

7.1 Requirements for Embedded ESG[13]

The specification defines the following two flavors for the embedded device:

1) **ESG as an extension of the eDVA eSAFE:** The eDVA eSAFE is updated as required to support the ESG functionality. This allows the re-use of existing E-DVA code base and BSS/OSS servers with minimal changes, potentially leading to fast product development and easy deployment.

2) **ESG as a separate eSAFE:** The ESG functionality is supported on a new eSAFE device defined specifically to support ESG functionality. This allows operators to use separate BSS/OSS servers than used for E-DVA. The ESG application can be developed independent of the E-DVA code base and can potentially support larger business voice customers than feasible with other embedded option.

When the eESG eSAFE is configured for dual-stack mode as specified in [PKT-EUE-PROV], the following requirements apply:

Note: In the following requirements, the term "IPv6 address" refers only to IPv6 Global Unicast Addresses (GUA), and does not include other IPv6 address types (e.g., does not include link-local addresses).

a) If the eESG has obtained WAN IP addresses from both address families (IPv4 and IPv6), then it MUST select a WAN IP address for its SIP signaling interface from the same address family (same IP version) as the primary IP address defined in [PKT-EUE-PROV].

Note: Even though a dual-stack eESG obtains WAN IP addresses from both address families, it operates as a single-stack device with-respect-to its WAN SIP signaling and management interfaces; i.e., these interfaces all use addresses within the same address family. The provisioning procedures defined in [PKT-EUE-PROV] provide a dual-stack eESG with a single IP address within each address family, in which case the SIP signaling and management interfaces will all use the same IP address. However, in the more general case where the device obtains multiple WAN IPv6 addresses, and IPv6 is the preferred IP version, the eESG is expected to use the algorithms defined in [RFC 3484] for address selection of the SIP signaling interface. In this case, the WAN IPv6

[12] Updated by ESG-N-11.0665-6, 11/1/16 by PO
[13] Updated by ESG-N-11.0665-6, 11/1/16 by PO

address selected for SIP signaling could be different than the WAN IPv6 address selected for management.

b) If the DNS IP address resolution of the P-CSCF FQDN provides both IPv4 and IPv6 addresses, the eESG MUST follow its normal P-CSCF IP address selection and failover procedures defined for initial registration in [PKT 24.229], with the exception that it limit P-CSCF address selection to only those addresses whose IP version match the IP version of the eESG SIP signaling WAN IP address.

c) The eESG MUST be capable of using both IP versions in the WAN media plane.

Note: A single 2-way media stream for an established 2-way call will be transported over a single IP version; either IPv4 or IPv6. However, in its role as a media-relay device for calls to/from the enterprise, the eESG needs to be able to support both IP versions in the WAN media plane simultaneously, to support the case where the calls terminate to a mix of IPv4-only and IPv6-only remote endpoints. The eESG is also required to support both media IP versions simultaneously within the scope of a single call. For example, for multi-party features such as call-transfer and 3-way calling, a single call can have two call legs with two remote endpoints that support different IP versions. Also, during early media sessions with a remote dual-stack endpoint, the eESG needs to be prepared to receive early media on either IP version.

d) During session establishment, the eESG MUST negotiate the IP address used for media from among the IPv4 and IPv6 addresses obtained during provisioning, as specified in Annex M of [PKT 24.229].

Note: [PKT 24.229] Annex K mandates the use of the "LITE" implementation of the Interactive Connectivity Establishment (ICE) protocol as defined in [RFC 5245] to negotiate the media IP address for single-stack IPv6 and dual-stack endpoints.

e) When the ICE-lite media IP address negotiation procedure yields multiple valid candidate pairs across both IP versions for a given media stream, the eESG designated as the "controlling agent" by the ICE-lite procedures MUST select a candidate pair whose IP version matches the locally configured preferred IP version for media defined in [PKT-EUE-DATA]. In addition, when selecting among multiple IPv6 candidate pairs (say when IPv6 is the preferred IP version), the "controlling agent" eESG MUST follow the address selection procedures mandated by ICE, and select the media WAN IP address based on the selection algorithms defined in [RFC 3484].

The architecture and requirements for eSAFE devices are defined in [eDOCSIS].

7.1.1 Embedded ESG as an Extension of eDVA eSAFE

This version of the embedded ESG extends the existing eDVA eSAFE to support the new ESG functionality, as shown in Figure 14. This version of the ESG is referred to as the "ESG E-DVA".

The ESG E-DVA MUST conform to the requirements in [PKT-RST-E-DVA] with the following additions:

1) The ESG E-DVA MUST support all the logical components in the ESG application and associated requirements as detailed in the Section 6 of this document.

2) The ESG E-DVA MUST support the additional ESG E-DVA provisioning requirements as defined in the Section 8 of this document.

3) The ESG E-DVA MUST support at least two physical customer-facing Ethernet ports, with one port dedicated to the broadband high-speed-data service, and the other port dedicated to Business Voice service.

4) The ESG E-DVA eDVA eSAFE MUST support at least 4 FXS ports.

The ESG E-DVA uses the same WAN IP address and the same device certificate that is assigned to the eDVA eSAFE, to communicate with the Service Provider network. Additionally, the eDVA eSAFE is assigned a LAN IP

address which the ESG uses to communicate with the enterprise SIP end points in the customer administered network. The configuration information for the ESG application is provided as part of the eDVA configuration file.

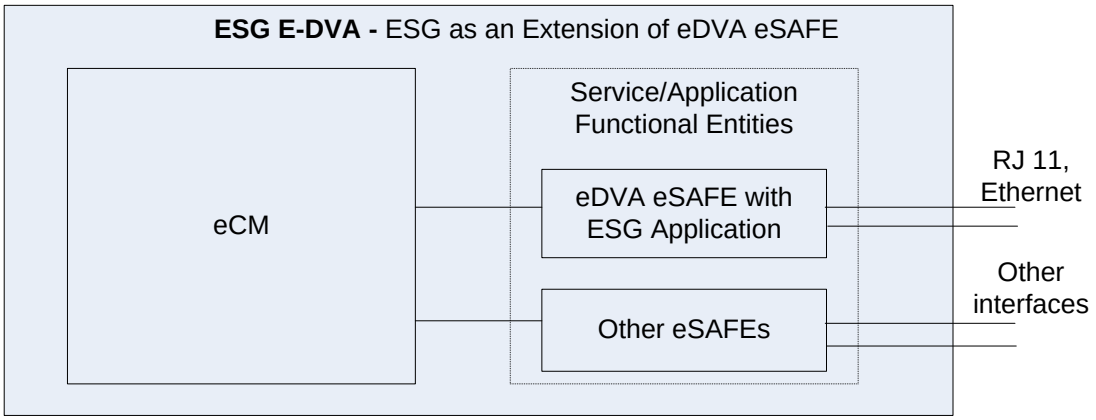

Figure 14 - Embedded ESG as an extension of eDVA eSAFE

7.1.2 Embedded ESG as a Separate eESG eSAFE[14]

This version of the embedded ESG supports the ESG functionality in a new eSAFE device called the eESG, as shown in Figure 15. This version of the ESG is referred to as the "E-ESG".

The E-ESG MUST support an eESG eSAFE as defined in this section. The E-ESG MUST support an eDVA eSAFE that conforms with the requirements in [PKT-RST-E-DVA]. The E-ESG eDVA eSAFE MUST support at least four FXS ports.

The eESG eSAFE must conform to the following requirements:

1) The eESG eSAFE MUST support all the logical components in the ESG application and associated requirements as detailed in the Section 6 of this document.

2) The eESG eSAFE MUST support the provisioning requirements as defined in the Section 8 of this document.

3) When the eESG eSAFE is configured for dual-stack mode as specified in [PKT-EUE-PROV], the following requirements apply:

Note: In the following requirements, the term "IPv6 address" refers only to IPv6 Global Unicast Addresses (GUA), and does not include other IPv6 address types (e.g., does not include link-local addresses).

a) If the eESG has obtained WAN IP addresses from both address families (IPv4 and IPv6), then it MUST select a WAN IP address for its SIP signaling interface from the same address family (same IP version) as the primary IP address defined in [PKT-EUE-PROV].

Note: Even though a dual-stack eESG obtains WAN IP addresses from both address families, it operates as a single-stack device with-respect-to its WAN SIP signaling and management interfaces; i.e., these interfaces all use addresses within the same address family. The provisioning procedures defined in [PKT-EUE-PROV] provide a dual-stack eESG with a single IP address within each address family, in which case the SIP signaling and management interfaces will all use the same IP address. However, in the more general case where the device obtains multiple WAN IPv6 addresses, and IPv6 is the preferred IP version, the eESG is expected to use the algorithms defined in [RFC 3484] for address selection of the SIP signaling interface. In this case, the WAN IPv6

[14] Updated by ESG-N-11.0665-6, 11/1/16 by PO

address selected for SIP signaling could be different than the WAN IPv6 address selected for management.

b) If the DNS IP address resolution of the P-CSCF FQDN provides both IPv4 and IPv6 addresses, the eESG MUST follow its normal P-CSCF IP address selection and failover procedures defined for initial registration in [PKT 24.229], with the exception that it limit P-CSCF address selection to only those addresses whose IP version match the IP version of the eESG SIP signaling WAN IP address.

c) The eESG MUST be capable of using both IP versions in the WAN media plane.

Note: A single 2-way media stream for an established 2-way call will be transported over a single IP version; either IPv4 or IPv6. However, in its role as a media-relay device for calls to/from the enterprise, the eESG needs to be able to support both IP versions in the WAN media plane simultaneously, to support the case where the calls terminate to a mix of IPv4-only and IPv6-only remote endpoints. The eESG is also required to support both media IP versions simultaneously within the scope of a single call. For example, for multi-party features such as call-transfer and 3-way calling, a single call can have two call legs with two remote endpoints that support different IP versions. Also, during early media sessions with a remote dual-stack endpoint, the eESG needs to be prepared to receive early media on either IP version.

d) During session establishment, the eESG MUST negotiate the IP address used for media from among the IPv4 and IPv6 addresses obtained during provisioning, as specified in Annex M [PKT 24.229]].

Note: [PKT 24.229] Annex K mandates the use of the "LITE" implementation of the Interactive Connectivity Establishment (ICE) protocol as defined in [RFC 5245] to negotiate the media IP address for single-stack IPv6 and dual-stack endpoints.

e) When the ICE-lite media IP address negotiation procedure yields multiple valid candidate pairs across both IP versions for a given media stream, the eESG designated as the "controlling agent" by the ICE-lite procedures MUST select a candidate pair whose IP version matches the locally configured preferred IP version for media defined in [PKT-EUE-DATA]. In addition, when selecting among multiple IPv6 candidate pairs (say when IPv6 is the preferred IP version), the "controlling agent" eESG MUST follow the address selection procedures mandated by ICE, and select the media WAN IP address based on the selection algorithms defined in [RFC 3484].

The E-ESG must conform to the following requirements:

1) The E-ESG MUST support at least two physical customer-facing Ethernet ports, with one port dedicated to the broadband high-speed-data service, and the other port dedicated to Business Voice service.

2) The E-ESG MUST support independent service states for the eDVA and eESG eSAFEs (e.g., where eESG is enabled and operational while eDVA is disabled).

The eESG eSAFE is assigned a different WAN IP address than assigned to the eDVA eSAFE. The eESG eSAFE is also assigned a LAN IP address which the ESG uses to communicate with the enterprise SIP endpoints in the customer administered network.

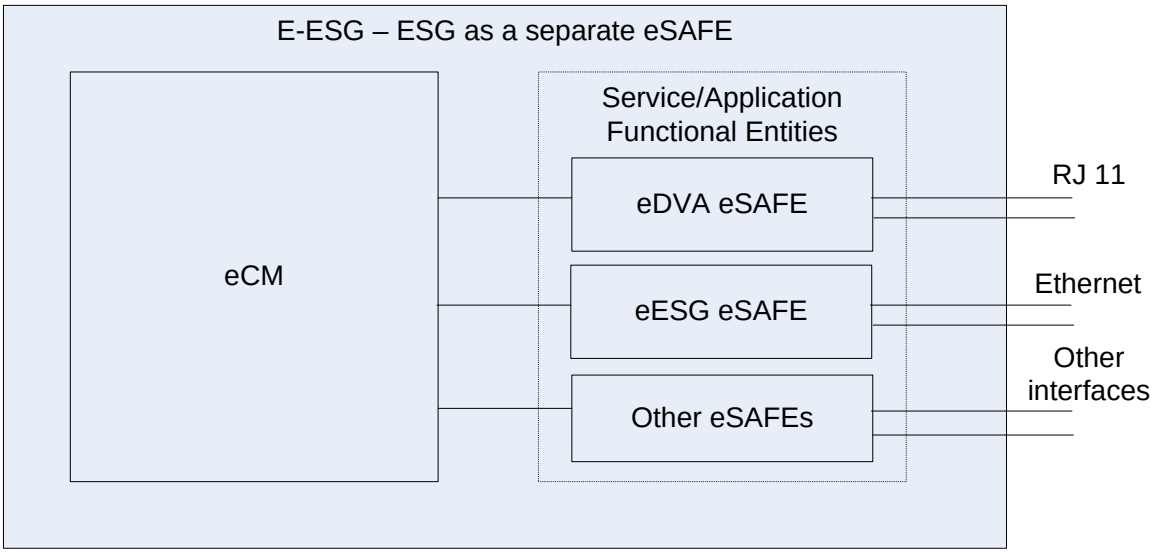

Figure 15 - Embedded ESG as a Separate eSAFE

7.2 Requirements for Stand-alone ESG[15]

This document refers to the stand-alone ESG as an "S-ESG" (see Figure 16). The S-ESG must conform to the following requirements:

1) The S-ESG MUST support at least one physical customer-facing Ethernet port dedicated to Business Voice service.
2) The S-ESG MUST support a single physical Ethernet port facing the Service Provider network.
3) The S-ESG MUST support all the logical components in the ESG application and associated requirements as detailed in the Section 6 of this document.
4) The S-ESG MUST support the provisioning requirements as defined in the Section 8 of this document.
5) When the S-ESG is configured for dual-stack mode as specified in Section 8.1.3, the following requirements apply:

Note: In the following requirements, the term "IPv6 address" refers only to IPv6 Global Unicast Addresses (GUA), and does not include other IPv6 address types (e.g., does not include link-local addresses).

a) If the S-ESG has obtained WAN IP addresses from both address families (IPv4 and IPv6), then it MUST select a WAN IP address for its SIP signaling interface that matches the locally configured IP version for signaling and management, as defined in Section 8.1.3,

Note: Even though a dual-stack S-ESG obtains WAN IP addresses from both address families, it operates as a single-stack device with-respect-to its WAN SIP signaling and management interfaces; i.e., these interfaces all use addresses within the same address family. Normally, a dual-stack S-ESG obtains only a single WAN IP address within each address family, in which case the SIP signaling and management interfaces all use the same IP address. However, in the more general case where the device obtains multiple WAN IPv6 addresses, and IPv6 is the preferred IP version, the S-ESG is expected to use the algorithms defined in [RFC 3484] for address selection of the SIP signaling interface. In this case, the WAN IPv6 address selected for SIP signaling could be different than the WAN IPv6 address selected for management.

[15] Updated by ESG-N-11.0665-6, 11/1/16 by PO

b) If the DNS IP address resolution of the P-CSCF FQDN provides both IPv4 and IPv6 addresses, the S-ESG MUST follow its normal P-CSCF IP address selection and failover procedures defined for initial registration in [PKT 24.229], with the exception that it limit P-CSCF address selection to only those addresses whose IP version match the IP version of the S-ESG SIP signaling WAN IP address.

c) The S-ESG MUST be capable of using both IP versions in the WAN media plane.

Note: A single 2-way media stream for an established 2-way call will be transported over a single IP version; either IPv4 or IPv6. However, in its role as a media-relay device for calls to/from the enterprise, the S-ESG needs to be able to support both IP versions in the WAN media plane simultaneously, to support the case where the calls terminate to a mix of IPv4-only and IPv6 only remote endpoints. The S-ESG is also required to support multiple media IP versions simultaneously within the scope of a single call. For example, for multi-party features such as call-transfer and 3-way calling, a single call can have two call legs with two remote endpoints that support different IP versions. Also, during early media sessions with a remote dual-stack endpoint, the S-ESG needs to be prepared to receive early media on either IP version.

d) During session establishment, the S-ESG MUST negotiate the IP address used for media from among the IPv4 and IPv6 addresses obtained during provisioning, as specified in Annex M of [PKT 24.229].

Note: [PKT 24.229] Annex K mandates the use of the "LITE" implementation of the Interactive Connectivity Establishment (ICE) protocol as defined in [RFC 5245] to negotiate the media IP address for single-stack IPv6 and dual-stack endpoints.

e) When the ICE-lite media IP address negotiation procedure yields multiple valid candidate pairs across both IP versions for a given media stream, the S-ESG designated as the "controlling agent" by the ICE-lite procedures MUST select an IP address whose version matches the locally configured preferred IP version for media defined in 8.1.3,. In addition, when selecting among multiple IPv6 candidate pairs (say when IPv6 is the preferred IP version), the "controlling agent" S-ESG MUST follow the address selection procedures mandated by ICE, and select the media WAN IP address based on the selection algorithms defined in [RFC 3484].

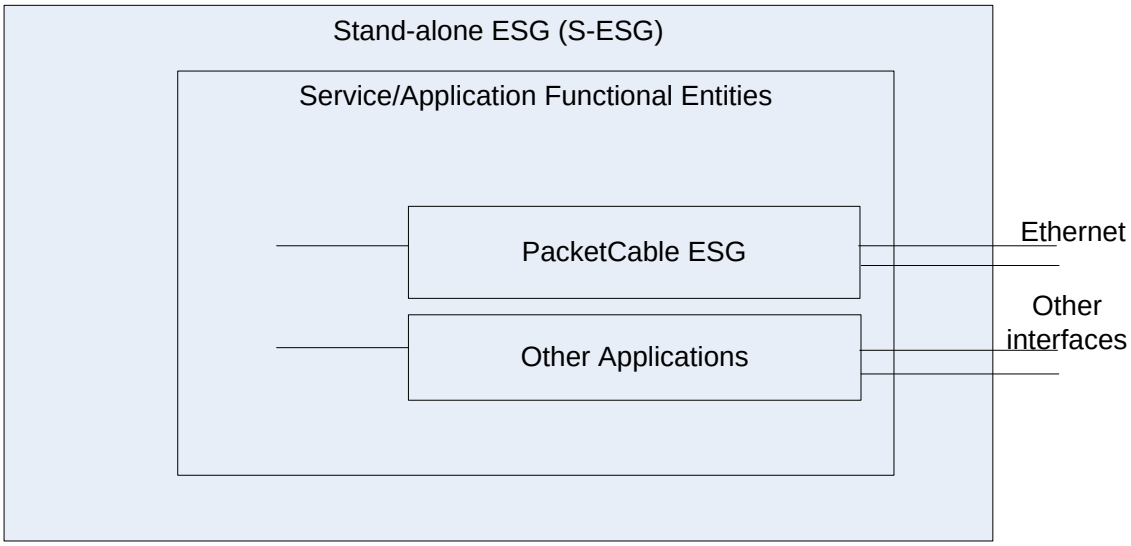

Figure 16 - Stand-alone ESG

8 OAM&P REQUIREMENTS

This section describes the normative requirements for Operations, Administration, Maintenance and Provisioning (OAM&P) functionality for ESG devices.

8.1 ESG Provisioning

This section describes the provisioning mechanisms used to configure and manage the PacketCable ESG. The configuration and management data model is defined in Annex A of this document.

PacketCable specifies a DHCP and SNMP based framework for clients that are embedded with DOCSIS cable modems and are not behind NAT and firewall devices (e.g., E-MTA, E-DVA). The details of this framework can be found in [PKT-PROV1.5], and [PKT-EUE-PROV].

PacketCable specifies an OMA DM based Provisioning & Management Framework for clients that are stand-alone and maybe behind NAT and firewall devices (e.g., UE). The details of this framework can be found in [PKT-UE-PROV].

The provisioning procedures defined in this specification are dependent on the type of ESG device (ESG E-DVA vs. E-ESG vs. S-ESG). The details of these procedures are provided in the sections below.

Note: Future releases of this specification may define additional ESG provisioning requirements.

8.1.1 ESG E-DVA Provisioning Requirements

As an extension of eDVA eSAFE, the ESG E-DVA MUST conform to the provisioning and management requirements in [PKT-RST-E-DVA]. The ESG E-DVA MUST support the management requirements defined in [PKT-RST-EUE-PROV], [PKT-EUE-DATA], and in Annex A of this document.

In addition, the eDVA component of the ESG E-DVA MUST report the following capability in DHCPv4 Option 60 and DHCPv6 Option 'CL_OPTION_MODEM_CAPABILITIES' in the DHCP messages:

Type	Length	Value	Comment
5.36	0	N/A	"Length=0" indicates that value is not relevant for this TLV

The ESG E-DVA, as explained in the device requirement Section 7.1.1, uses one eSAFE for both RST and ESG applications. As a result, a reset operation on this eSAFE will reset both ESG and RST applications. The ESG E-DVA is not required to support reset of RST and ESG applications individually. Additionally, the ESG E-DVA is not required to maintain its service on the enterprise side if the eCM becomes non-operational for any reason.

The ESG E-DVA MUST report the corresponding value of 'Service Interruption Impact' as defined in [PKT-EUE-PROV]. In doing so, the ESG E-DVA MUST support the following additional requirements for ESG application:

- The ESG application MUST indicate a value of 'significant(1)' if it is configured for services and in an "operational" state.
- The ESG application MUST indicate value 'none(2)' to ESG E-DVA in all other cases.
- If ESG E-DVA reports the value of 'significant(1)' indicated by the ESG application, it MUST also report the reason in the `esafeDevServiceIntImpactInfo` MIB Object.
- In case, both applications RST and ESG are indicating the value 'significant(1)', the ESG E-DVA MUST report in the `esafeDevServiceIntImpactInfo` MIB Object reasons for both applications.

8.1.2 E-ESG Provisioning Requirements

The eCM component in the E-ESG MUST conform to the requirements in DOCSIS, [PKT-RST-E-DVA], and [PKT-EUE-PROV] specifications.

The eDVA component, in the E-ESG, MUST conform to the provisioning and management requirement in [PKT-RST-E-DVA].

Also, the eESG component in the E-ESG MUST conform to the requirements of the eSAFE device as described in [eDOCSIS] specification.

The eESG component, in the E-ESG, MUST conform to the following requirements (the 'pkt-eue-prov-x' reference points are defined in [PKT-EUE-PROV]):

- The eESG MUST implement a DHCP Client. The eESG MUST NOT use DHCP if it is preconfigured with an IP address, DNS Server, and configuration server address and configuration file name. Otherwise, the eESG SHOULD use DHCP to identify itself to the Service Provider network. DHCP servers use this information to provide IP configuration information such as IP address, DNS server address, configuration server address and name of the configuration file for the eESG. If using DHCP to obtain IP configuration, the eESG DHCP Client MUST conform to the pkt-eue-prov-1 interface requirements applicable for the Basic Provisioning Flow as defined in [PKT-EUE-PROV] with following considerations:
- The eESG MUST use a value of "EESG" in the DHCP Option 43, sub-option 2 for IPv4.
- The eESG MUST use a value of "EESG" in the CL_OPTION_DEVICE_TYPE(2) for IPv6.
- The eESG is not required to include DHCPv4 Option 60 and DHCPv6 Option 'CL_OPTION_MODEM_CAPABILITIES' in the DHCP messages. If an eESG decides to use Option 60 in the DHCP messages, then care should be taken to only advertise the functionality supported by the eESG. This specification does not define standard TLVs for eESG capabilities.
- The eESG MUST perform a protocol (e.g., HTTP, TFTP) exchange to download its configuration file. The 'siaddr' and 'file' fields of the DHCP ACK are used to locate the configuration file. Specific details of the protocol exchange and the content of the configuration file are out of scope of this version of the specification.
- The eESG MUST implement and use the pkt-eue-prov-2 interface.
- The eESG MUST NOT use the pkt-eue-prov-3 interface.
- The eESG MAY implement and use the pkt-eue-prov-4 interface. The eESG MUST implement and use pkt-eue-prov-4 if BASIC.2 Provisioning Flow is requested.
- The eESG MAY implement and use the pkt-eue-prov-5 interface. The eESG SHOULD implement an interface to support the download of the ESG configuration file.
- The eESG MUST implement and use the pkt-eue-prov-6 interface.
- The eESG MUST have its own MAC address, different from the MAC address of the eCM and eDVA.
- The eESG MUST have its own WAN IP address, different from the IP address(es) of the eCM and eDVA.
- The eESG MUST be able to operate in environments where the eESG WAN IP address is in the same IP subnet, or in a different IP subnet, as the eCM and eDVA.
- The eESG MUST start the provisioning process immediately after the eCM component is in the "operational" state.
- Once the eESG is in-service and operational, it MUST maintain its service on the enterprise side if the eCM becomes non-operational for any reason. This would enable the eESG to reject incoming call requests from the enterprise with an error code or announcement.

If the eESG supports pkt-eue-prov-4 (i.e., SNMP) interface, it MUST indicate the service status in the esafeDevServiceIntImpact MIB Object according to the following logic.

- The eESG MUST report value 'significant(1)' if the eESG is configured for services and is in "operational" state.
- The eESG MUST report value 'none(2)' on all other cases.
- The eESG MUST indicate the particular reason of its inability to provide the configured services in the `esafeDevServiceIntImpactInfo` MIB Object.

8.1.3 S-ESG Provisioning Requirements[16]

The S-ESG MUST implement a DHCP Client. The S-ESG MUST NOT use DHCP if it is preconfigured with an IP address, DNS Server, and configuration server address and configuration file name. The S-ESG SHOULD use DHCP to identify itself to the Service Provider network. DHCP servers use this information to provide WAN IP configuration information such as IP address, DNS server address, configuration server address and name of the ESG configuration file to the S-ESG. If using DHCP to obtain IP configuration, the S-ESG DHCP Client must conform to the pkt-eue-prov-1 interface requirements applicable for the Basic Provisioning Flow [PKT-EUE-PROV] with following considerations:

- The S-ESG MUST use a value of "SESG" in the DHCP Option 43, sub-option 2 for IPv4.
- The S-ESG MUST use a value of "SESG" in the CL_OPTION_DEVICE_TYPE(2) for IPv6.
- The S-ESG is not required to include DHCPv4 Option 60 or DHCPv6 Option 'CL_OPTION_MODEM_CAPABILITIES' in the DHCP messages. If an S-ESG decides to use either of these options in the DHCP messages, care should be taken to only advertise the functionality supported by the S-ESG. This specification does not define standard TLVs for S-ESG capabilities.
- If DHCP is used to acquire the IP address, the S-ESG MUST do so according to the logic of [RFC 2131] for IPv4 and [RFC 3315] for IPv6.

The S-ESG MUST support a configuration option to control whether it is operating in single-stack or dual-stack mode, where the single-stack mode supports only one IP version (IPv4 or IPv6), while the dual-stack mode supports both IP versions (IPv4 and IPv6) on the WAN interface. When configured for dual-stack mode, the S-ESG MUST support configuration options that identify the following:

- The IP version of the WAN SIP signaling and management interfaces.
- The preferred IP version in the WAN media plane; i.e., the IP version that is preferred when both versions are supported by the remote endpoint.

The provisioning mechanisms and attribute names and values associated with these configuration options are outside the scope of this specification.

The S-ESG MUST perform a protocol (e.g., HTTP, TFTP) exchange to download its configuration file. If the DHCP is used, the corresponding DHCP options are used to locate the configuration file ('siaddr' and 'file' fields for DHCPv4 and CL_OPTION_TFTP_SERVERS and CL_OPTION_CONFIG_FILE_NAME for DHCPv6 as defined in [CL-CANN-DHCP-REG]). Specific details of the protocol exchange and the content of the configuration file are out of scope of this version of the specification.

The S-ESG can begin the DHCP process as soon at the device is switched on and the physical interface for the WAN side network is enabled.

The S-ESG, must conform to the following requirements (the 'pkt-eue-prov-x' reference points listed below are defined in [PKT-EUE-PROV]).

[16] Updated by ESG-N-11.0665-6, 11/1/16 by PO

- The S-ESG SHOULD implement and use the pkt-eue-prov-2 interface.
- The S-ESG MUST NOT use the pkt-eue-prov-3 interface.
- The S-ESG MAY implement and use the pkt-eue-prov-4 interface.
- The S-ESG MAY implement and use the pkt-eue-prov-5 interface.
- The S-ESG SHOULD implement an interface like pkt-eue-prov-5 to support the download of the ESG configuration file.
- The S-ESG SHOULD implement and use the pkt-eue-prov-6 interface.

8.2 ESG Provisioning Additional Features

8.2.1 Persistent Configuration Support

The configuration of the ESG by the Operator involves numerous data objects and can be quite complicated in making sure that the configuration data is consistent and provides the desirable set of functionality. Creating configuration file by the PacketCable Provisioning System each time the ESG goes through the reset, and then downloading this file to ESG may become time consuming and challenging task for mass deployments.

To address this issue, the S-ESG, and E-ESG with dedicated eSAFE MUST store the entire configuration data in the configuration file in the Non-Volatile storage.

The E-ESG with shared eSAFE MAY store the entire configuration data in the Non-Volatile storage. If the ESG stores the configuration data to Non-Volatile storage, it MUST also store the SHA-1 hash value of the last known downloaded configuration file which carried this configuration data. The ESG MUST calculate the SHA-1 hash value as required in [PKT-EUE-PROV].

8.2.2 Battery Support

An ESG E-DVA supporting Battery Backup MUST support the requirements specified in the Battery Backup MIB Specification [CL-SP-MIB-BB].

8.3 Provisioning Application Level Gateway[17]

This section contains the normative requirements for the two types of Provisioning ALGs described in Section 5.2.3; the CWMP ALG and the HTTP ALG.

8.3.1 CWMP ALG

The ESG SHOULD support the CWMP ALG function. If the ESG supports the CWMP ALG function, it MUST comply with the CWMP ALG requirements specified in this document.

As described in Section 5.2.3.1, the CWMP ALG provides both IP-layer and application-layer interworking for CWMP traffic between TR-069-compliant ESEs and their ACS.

[17] New sections added by ESG-N-12.0691-8 on 11/4/16 by PO

8.3.1.1 Mapping CWMP ALG to the TR-069 Architecture

In general, the CWMP ALG does not map to a specific entity within the TR-069 architecture. Rather, the ALG sits on the Service Provider/Enterprise WAN/LAN boundary and provides a NAT-traversal solution that enables transparent routing of CWMP messages between the ESEs and their ACS.

Although the CWMP ALG doesn't fully realize any single TR-069 entity, it does leverage the following two capabilities of the Gateway entity defined in Annex F of [TR-069].

1. TR-069 Annex F describes the topology option where a single Gateway provides network access to multiple CPE devices located on a LAN, and where the ACS manages both the Gateway and the CPEs. TR-069 provides a mechanism for the ACS to learn the relationship between the Gateway and its subordinate CPEs. Basically, the CPE learns the Gateway Identity during DHCP discovery, and then reports that Gateway Identity to the ACS in the CWMP Inform request. For the ESG case, the ESE and CWMP ALG will not be required to support the TR-069 mechanism where the ESE learns the Gateway Identity. Instead, the CWMP ALG will add its Gateway Identity to each Inform request received from an ESE before forwarding the request to the ACS. This will enable the ACS to resync its stored ConnectionRequestURL entries when the WAN IP address of the ESG changes, as described in section 8.3.1.8.

2. TR-069 Annex F also defines a Gateway data model containing a ManageableDevice table. The Gateway populates this table with the Device Identities contained in the Inform requests received from CPE devices. The CWMP ALG will also support this table, populating it with the Device Identities contained in the Inform requests received from the ESEs. This ESE Device Identity information will enable the CWMP ALG to release stale entries in its address mapping table as described in section 8.3.1.5.

   Note: Support of the above TR-069 Gateway capabilities are not normatively mandated in this version of the document. Normative requirements for these capabilities will be added to a future version of the specification.

8.3.1.2 CWMP ALG Support of Traditional NAT

The CWMP ALG MUST support the traditional NAT functionality specified in [RFC 2663] to transparently route TCP and UDP CWMP messages between an Enterprise SIP endpoint in the Enterprise network and an ACS in the Service Provider network. The CWMP ALG SHOULD support the traditional NAT requirements specified in [RFC 4787].

8.3.1.3 General NAT Traversal for any CWMP Request

One of the primary functions of the CWMP ALG is to enable NAT traversal at both the IP layer and the application layer for CWMP request/response transactions between an ESE and its ACS.

To accomplish this, the CWMP ALG MUST support two configurable data attributes:

- an ACS-LAN-URL that identifies the LAN IP address:port on the ESG that is allocated to receive CWMP requests from ESEs served by this CWMP ALG, and
- an ACS-WAN-URL that identifies the publically routable URL of the ACS associated with this CWMP ALG.

On receiving a CWMP request on the LAN IP address:port identified by the configured ACS-LAN-URL, the CWMP ALG MUST update the Request URI to the configured ACS-WAN-URL, and update the destination IP address:port of the request with the IP address:port associated with the ACS-WAN-URL, and forward the request to the ACS.

8.3.1.4 NAT Traversal for the CWMP Inform Request

Special application-level NAT traversal processing is required for the CWMP Inform request to interwork the IP address information carried in the ConnectionRequestURL. Basically, the CWMP ALG has to assign a WAN

ConnectionRequestURL for each incoming LAN ConnectionRequestURL received from an ESE, and maintain a WAN-to-LAN ConnectionRequestURL address mapping table to route subsequent unsolicited connection requests from the ACS to the correct ESE.

On receiving a CWMP Inform request on its configured ACS-LAN-URL from an ESE device that does not have an entry in the WAN-to-LAN ConnectionRequestURL address mapping, the CWMP ALG MUST perform the following:

a) assign a new WAN ConnectionRequestURL for the ESE that identifies or resolves to the ESG WAN IP address:port/path to receive subsequent unsolicited connection requests from the ACS,

b) update the Inform ConnectionRequestURL with the newly assigned WAN ConnectionRequestURL,

c) add an entry to the address mapping table that contains the WAN ConnectionRequestURL, the mapped LAN ConnectionRequestURL, and the Device Identity of the ESE that sent the Inform request, and

d) forward the request to the ACS as specified in Section 8.3.1.3.

Note: Other than the requirement that the WAN ConnectionRequestURL can be uniquely mapped to the associated LAN ConnectionRequestURL, the details for assigning/building a new WAN ConnectionRequestURL are left up to vendor implementation. The WAN ConnectionRequestURL could be made unique by assigning a different WAN port per URL. Or, all WAN ConnectionRequestURLs could be assigned to the same WAN port, and individually made unique by being given a unique path.

On receiving a CWMP Inform request on its configured ACS-LAN-URL from an ESE device that has an entry in the ConnectionRequestURL address mapping table, and the ConnectionRequestURL in the Inform matches the LAN ConnectionRequestURL for that device in the address mapping table, the CWMP ALG MUST perform the following:

a) update the Inform ConnectionRequestURL with the device's WAN ConnectionRequestURL specified in the table, and

b) forward the request to the ACS as described in Section 8.3.1.3.

On receiving a CWMP Inform request on its configured ACS-LAN-URL from an ESE device that has an entry in the ConnectionRequestURL address mapping table, and the ConnectionRequestURL in the Inform does not match the LAN ConnectionRequestURL for that device in the address mapping table, the CWMP ALG MUST perform the following:

a) update the LAN ConnectionRequestURL in the address mapping table entry for that device to the new ConnectionRequestURL identified in the received Inform request,

b) update the Inform ConnectionRequestURL with the device's WAN ConnectionRequestURL specified in the table, and

c) forward the request to the ACS as described in Section 8.3.1.3.

Note: The above requirement mandates that for the case where an ESE with an entry in the address mapping table sends an Inform request containing a ConnectionRequestURL that is different than the LAN ConnectionRequestURL in the ESE's table entry, the CWMP ALG reuses the existing WAN ConnectionRequestURL already assigned to the device instead of assigning a new WAN ConnectionRequestURL.

The CWMP ALG SHOULD provide a mechanism to rate-limit Inform requests.

8.3.1.5 Releasing Stale ConnectionRequestURL Address Mapping Entries

The CWMP ALG SHOULD implement an algorithm to release the WAN-to-LAN ConnectionRequestURL address mapping entries when the ESE that established the mapping is no longer available (say the ESE is disconnected from the Enterprise network).

Note: The algorithm for releasing this mapping data is not specified, but could consist of one or more of the following mechanisms:

- The CWMP ALG detects lack of CWMP activity to/from the ESE for an extended period (this could include lack of activity for other protocols such as SIP). This mechanism would have to take into account operator policy that governs whether and how often there is periodic CWMP communication between an ESE and its ACS.
- The CWMP detects that the LAN IP address of an address mapping entry is reused by a new ESE device (i.e., if an ESE sends an Inform containing a ConnectionRequestURL that has the same LAN IP address as an existing address mapping entry for another ESE device, then remove that conflicting entry).
- The CWMP ALG detects that the ESE is no longer accessible over IP. Note that status change in the local ARP table cannot be used to imply IP connectivity since ARP entries tend to time out quickly, and ARP entries don't exist for ESEs located on a different subnet than the ESE LAN subnet. If a periodic ICMP ping is used, then it may detect a loss of connectivity due to a momentary network outage or overload, and not because the ESE has been removed from the network. A ping also has the down-side that it would respond positively for the case where the ESE was removed from the network and a non-TR-069 device obtained the same IP address.
- The CWMP ALG releases the oldest inactive ESE entries once memory resources associated with the mapping data exceed some threshold.

8.3.1.6 Routing Unsolicited ACS Connection Requests

On receiving a valid ACS connection request on a WAN IP address:port associated with the WAN ConnectionRequestURL of an entry in the address mapping table, the CWMP ALG MUST update the received Request URI of the request to the mapped LAN ConnectionRequestURL, and forward the request to the target ESE. The CWMP ALG MUST provide a mechanism to rate-limit these requests.

Note: An ACS connection request is considered valid if it complies with the ACS connection request requirements specified in [TR-069], and if the Request URI points to an entry in the WAN-to-LAN ConnectionRequestURL address mapping table.

The CWMP ALG MUST ignore all other requests received on any WAN IP address:port associated with the WAN ConnectionRequestURLs of entries in the address mapping table, and all ACS connection requests that comply with TR-069 but that do not match a WAN ConnectionRequestURL in the address mapping table.

8.3.1.7 ACS Redirect

The CWMP ALG MUST interwork HTTP redirects received from the ACS such that the redirected URL provided to the ESE resolves to the LAN interface of the CWMP ALG.

Note: Providing the ESE with a redirected URL that resolves to the CWMP ALG LAN interface ensures that the subsequent CWMP requests sent by the ESE to the redirected URL will traverse the CWMP ALG.

8.3.1.8 Resyncing State after ESG Outage

The CWMP ALG MUST store the WAN to LAN ConnectionRequestURL address mapping table in non-volatile memory so that it can survive a restart or power-cycle of the ESG device. In addition, the CWMP ALG MUST provide a management mechanism to clear this mapping information (say, when an ESG is to be re-deployed in a new location).

If the ESG obtains a new WAN IP address, then the CWMP ALG MUST update the WAN ConnectionRequestURL entries of the address mapping table to align with the new IP address.

Note: It is assumed that the ACS performs a similar function to synchronize its ConnectionRequestURL entries associated with the ESEs subordinate to this CWMP ALG. Specifically, it is assumed that the ACS is configured with a policy such that whenever it receives notification from the CWMP ALG that the ALG's IP address has changed, the ACS updates the ConnectionRequesURI of each ESE subordinate to the CWMP ALG to align with the new IP address.

In order to enable the ACS to learn the mapping between the CWMP ALG and its subordinate ESEs, the CWMP ALG will include its Gateway Identity in each CWMP Inform request relayed to the ACS from an ESE.

Note: The Gateway Identity attribute referred to above is described in Annex F of [TR-069]. In this case, the CWMP ACS appears as a Gateway to the ACS. However, the CWMP ACS does not need to perform all the functions of the Gateway as specified in TR-069 Annex F. Rather, the CWMP ALG simply adds its own identity to the CWMP Inform request received from its subordinate ESEs before forwarding the Inform request to the ACS.

Note: The above informative text regarding inclusion of the Gateway Identity in the Inform request will be changed to a normative requirement in a future version of this specification.

8.3.1.9 CWMP ALG Security

8.3.1.9.1 TLS

The CWMP ALG MUST support TLS server procedures on its LAN interface and TLS client procedures on its WAN interface, following the TLS requirements described in [TR-069].

When its ACS-LAN-URL is configured with an HTTPS URL, the CWMP ALG MUST:

a) act as a TLS server, and accept TLS connection requests from the ESE, and

b) accept CWMP requests from the ESE only over the established TLS connection.

When its ACS-WAN-URL is configured with an HTTPS URL, the CWMP ALG MUST:

a) act as a TLS client, and establish a TLS session with the ACS for each CWMP connection request received from the ESE, and

b) send CWMP requests to the ACS over the established TLS sessions.

The CWMP ALG MUST support the case where both the ACS-WAN-URL and the ACS-LAN-URL have the same scheme (both URLs are HTTP, or both are HTTPS). The CWMP ALG SHOULD support the case where the ACS-WAN-URL and the ACS-LAN-URL have different schemes (one HTTP and the other HTTPS).

8.3.1.9.2 Digest Authentication

The CWMP ALG plays no role in the Digest authentication procedures between the ESE and ACS. However, since the CWMP ALG modifies the payload of CWMP messages, the ESE and ACS cannot use the Digest authentication "qop" option of "auth-int", defined in [RFC 2617].

8.3.2 HTTP ALG

The ESG SHOULD support the HTTP ALG function. If this function is supported, it MUST comply with the HTTP ALG requirements specified in this document.

As described in Section 5.2.3.2, the HTTP ALG provides both IP-layer and application-layer interworking for HTTP traffic between an ESE and its provisioning server. The application-layer interworking consists only of interworking HTTP Request-URIs and redirected URLs between the LAN and WAN networks.

8.3.2.1 HTTP ALG Support of Traditional NAT

The HTTP ALG MUST support the traditional NAT functionality specified in [RFC 2663] to transparently route TCP and UDP packets between an Enterprise SIP endpoint in the Enterprise network and a provisioning server in the Service Provider network. The HTTP ALG SHOULD support the traditional NAT requirements specified in [RFC 4787].

8.3.2.2 General NAT Traversal for any HTTP Request

The HTTP ALG MUST support two configurable data attributes:

- a Provisioning-Server-LAN-URL that identifies the LAN IP address:port on the ESG that is allocated to receive HTTP provisioning requests from ESEs served by this HTTP ALG, and
- a Provisioning-Server-WAN-URL that identifies the publically routable URL of the Provisioning Server associated with this HTTP ALG.

On receiving an HTTP request on the LAN IP address:port identified by the configured Provisioning-Server-LAN-URL, the HTTP ALG MUST update the Request URI to the configured Provisioning-Server-WAN-URL, and forward the request to the Provisioning Server.

8.3.2.3 HTTP Redirect

The HTTP ALG MUST interwork HTTP redirects received from the provisioning server such that the redirected URL provided to the ESE resolves to the LAN interface of the HTTP ALG.

Note: Providing the ESE with a redirected URL that resolves to the HTTP ALG LAN interface ensures that the subsequent HTTP requests sent by the ESE to the redirected URL will traverse the HTTP ALG.

8.3.2.4 HTTP ALG Security

8.3.2.4.1 TLS

The HTTP ALG MUST support TLS server procedures on its LAN interface and TLS client procedures on its WAN interface.

When its Provisioning-Server-LAN-URL is configured with an HTTPS URL, the HTTP ALG MUST:

a) act as a TLS server, and accept TLS connection requests from the ESE, and

b) accept HTTP requests from the ESE only over the established TLS connection.

When its Provisioning-Server-WAN-URL is configured with an HTTPS URL, the HTTP ALG MUST:

a) act as a TLS client, and establish a TLS session with the provisioning server when it receives an indication that the ESE wants to contact the provisioning server.

Note: The indication that the ESE wants to contact the provisioning server could take various forms; e.g., a TCP connection request or a UDP HTTP request from the ESE.

b) Send HTTP requests to the provisioning server over the established TLS session.

The HTTP ALG MUST support the case where both the Provisioning-Server-WAN-URL and the Provisioning-Server-LAN-URL have the same scheme (both URLs are HTTP, or both are HTTPS). The HTTP ALG SHOULD

support the case where the Provisioning-Server-WAN-URL and the Provisioning-Server-LAN-URL have different schemes (one HTTP and the other HTTPS).

8.3.2.4.2 HTTP Digest Authentication

The HTTP ALG plays no role in the HTTP Digest authentication procedures between the ESE and the provisioning server. However, since the HTTP ALG modifies the payload of the HTTP messages, the ESE and provisioning server cannot use the Digest authentication "qop" option of "auth-int", defined in [RFC 2617].

9 SECURITY REQUIREMENTS[18]

The ESG is required to provide security only on the SIP signaling interface to the Service Provider network (pkt-esg-sig-1). The security requirements for all other reference points are not specified in this document.

The ESG security requirements for pkt-esg-sig-1 depend on the type of business voice service being supported, as follows:

- For Hosted IP Centrex service, the ESG MUST support the SIP Digest authentication and TLS security requirements defined for the UE in [PKT 24.229].
- For SIP Trunking Service, the ESG MUST support the security requirements defined for the SIP-PBX in [SIPconnect1.1].

In addition, the ESG MUST support a configuration option that enables SIP Digest to be enabled or disabled.

When operating in "Registration Mode", the ESG MUST support the ability to initiate a TLS session at the time of SIP-PBX registration using the procedures defined in [RFC 3329].

In addition, the ESG MUST support an option to initiate the TLS session at ESG startup time as defined in [SIPconnect1.1].

The ESG MUST provide configuration data to govern behavior related to the selection between the procedures in [SIPconnect1.1] and [RFC 3329].

When providing SIP Trunking Service and operating in "Registration Mode", the ESG MUST support a configuration option that allows SIP Digest authentication to be enabled or disabled. If the option is set to "disabled", then the SBC component of the ESG passes 401/407 challenges on to the connected SIP-PBX. If the option is set to "enabled", the ESG MUST support SIP Digest authentication in order to answer authentication challenges issued by the SP-SSE (PacketCable 2.0 network), particularly in the context of SIP-PBX registration.

When providing SIP Trunking Service and operating in either "Registration" or "Static" mode, the ESG MUST support the TLS Mutual Authentication model.

When TLS is enabled, the ESG MUST support the PKI as defined in [PKT1.5-SEC] Specification.

[18] Updated by ESG-N-12.0682-8 on 10/28/16 by PO

Annex A ESG Object Model

A.1 ESG Object Model Overview

The UML (Unified Modeling Language) specification [ISO/IEC 19501] is used to define the Object Model for the ESG functions specified in this document. This Object Model organizes the ESG data items identified in Sections 6 through 9 into logical objects and attributes, and describes the relationships among these objects. The Object Model is generic in the sense that it provides a single general data definition that can be used to generate data files of different forms; say MIB-based, or XML-based.

A.1.1 SBC Function Use Cases

The ESG object model must be sufficiently flexible to accommodate the wide range of forms that the SBC can take; whether it's operating only as a SIP-aware NAT, or if it also plays the role of a SIP proxy or a SIP B2BUA.

A.1.1.1 SBC as SIP-aware NAT[19]

In its simplest form, the SBC serves a single SIP-PBX or SIP endpoint that is fully compatible with PacketCable 2.0. In this case the SBC operates as a SIP-aware NAT that provides IP address interworking between the Enterprise and PC 2.0 address space, but is otherwise transparent to SIP signaling. Figure 17 shows the ESG objects required to support this minimal example, where the ESG contains a single SBC object that in turn contains a single pair of ESE objects representing the single Enterprise SIP Entity. (ESG support of multiple ESEs is described in Section A.1.1.3.) The ESE(wan) object contains the public address information of the Enterprise SIP Entity, while the ESE(lan) object contains the private address information of the Enterprise SIP Entity.

The ESE(wan) object contains two public addresses:

- ESE-name(wan) is an alpha-numeric string that identifies the publically known identity of the ESE. It can be a SIP-URI which is the AOR assigned by the Service Provider to a registering SIP-PBX or SIP endpoint. Or, it can be the FQDN assigned by the Service Provider to a static mode SIP-PBX.
- ESE-location(wan) contains an IP address:port of the ESE within the public address space of the PC 2.0 network. It is the registered contact address of a registering SIP-PBX or SIP endpoint. Or, it is the IP address:port associated with the FQDN of a static-mode SIP-PBX.

Normally, the ESE(wan) object contains both the name and location address attributes. However, there are cases where the ESE(wan) contains only one of these two public addresses. For example, the Service Provider may choose to configure the routing data such that the PC 2.0 network identifies a static mode SIP-PBX using its public IP address:port, in which case there is no need for an ESE-name(wan) identifying the ESE FQDN. This will be described in more detail later in this section.

The ESE(lan) object contains the private address equivalents of the public addresses contained in ESE(wan).

- ESE-name(lan) is an alpha-numeric string that identifies the AOR assigned by the Enterprise to the SIP-PBX or SIP endpoint.
- ESE-location(lan) contains the IP address:port of the ESE within the Enterprise network. It is either statically configured by the Enterprise on the ESG, or discovered by the ESG via the registration procedure (via [1] REGISTER shown in Figure 17).

For Enterprise SIP Endpoints that register, the ESE(lan) object essentially acts as a registrar in the LAN network, maintaining the binding between the ESE AOR (lan) and the ESE contact address (lan).

[19] Updated by ESG-N-12.0682-8 on 10/28/16 by PO

The ESE(lan) also contains an attribute called "Interworking-rule-set→" that identifies a file containing the header manipulation rules that must be applied to achieve interworking between the Enterprise SIP Entity and the PacketCable 2.0 network. In this example the rule-set is empty, since the ESE is compatible with PC 2.0.

In addition to the ESE object pair, the SBC contains administrative and operational state attributes.

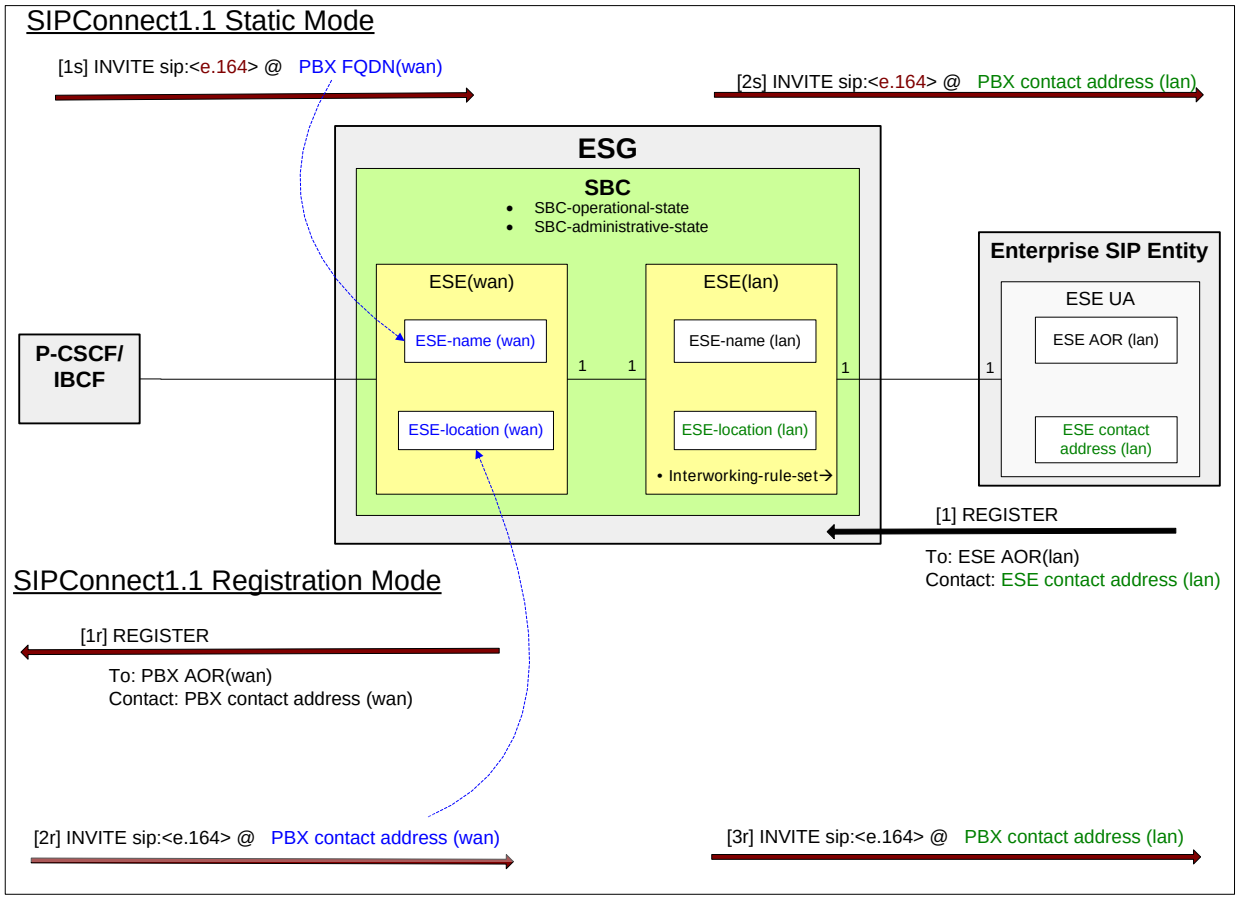

Figure 17 - Simple Pass-Thru ESG

Figure 17 also shows how the ESE(wan) and ESE(lan) objects are used to route SIP requests for the case where the Enterprise SIP Entity is a SIP-PBX.

- If the ESG is operating in the registration mode, then request (1r) conveys the public location of the SIP-PBX to the PC 2.0 network (the public location is usually specified in the form of an IP address:port). Once the ESE is registered with the PC 2.0 network, the PC 2.0 network can initiate a DID call by sending request [2r] where the public location of the target SIP-PBX is identified in the host-name of the INVITE Request-URI. On receiving [2r], the SBC identifies the target ESE(wan) by matching the Request-URI host-port to the ESE-location(wan). In this example, where the SIP-PBX supports [SIPconnect1.1], the SBC modifies the host-name in the Request-URI of [3r] to the registered LAN contact address of the target SIP-PBX.
- If the ESG is operating in the static mode, then the PC 2.0 network initiates a DID call by sending request [1s] where the FQDN of the target SIP-PBX is identified in the host-name of the INVITE Request-URI. On receiving [1s], the SBC identifies the target ESE(wan) by matching the Request-URI host-port to the ESE-name(wan).As in the previous case, the ESG modifies the host-name in the Request-URI of [2s] to the registered LAN contact address of the target SIP-PBX.

A.1.1.2 Adding Support for B2BUA

Section A.1.1.1 describes the simple case where the SBC is operating as a SIP-aware NAT, but is otherwise transparent to SIP signaling. If the SBC is also operating as a B2BUA, then the object model is extended such that the ESE(wan) becomes the PC 2.0 network-facing SIP User Agent, while the ESE(lan) becomes the Enterprise-facing SIP User Agent. To support this case, data attributes required by a SIP User Agent such as SIP transaction timers and next-hop address are added to ESE(wan). Certain SIP UA attributes such as Private Identity and Digest Credentials are required only for hosted IP-Centrex or registration-mode SIP Trunking service, where the ESE(wan) must support PacketCable 2.0 registration procedures. These additional attributes are not required for static-mode SIP Trunking service.

Similar SIP UA data attributes are required by the ESE(lan); however, these are considered internal to the ESG and are not reflected in the data model. The updated ESG data model to support the B2BUA case is shown in Figure 18.

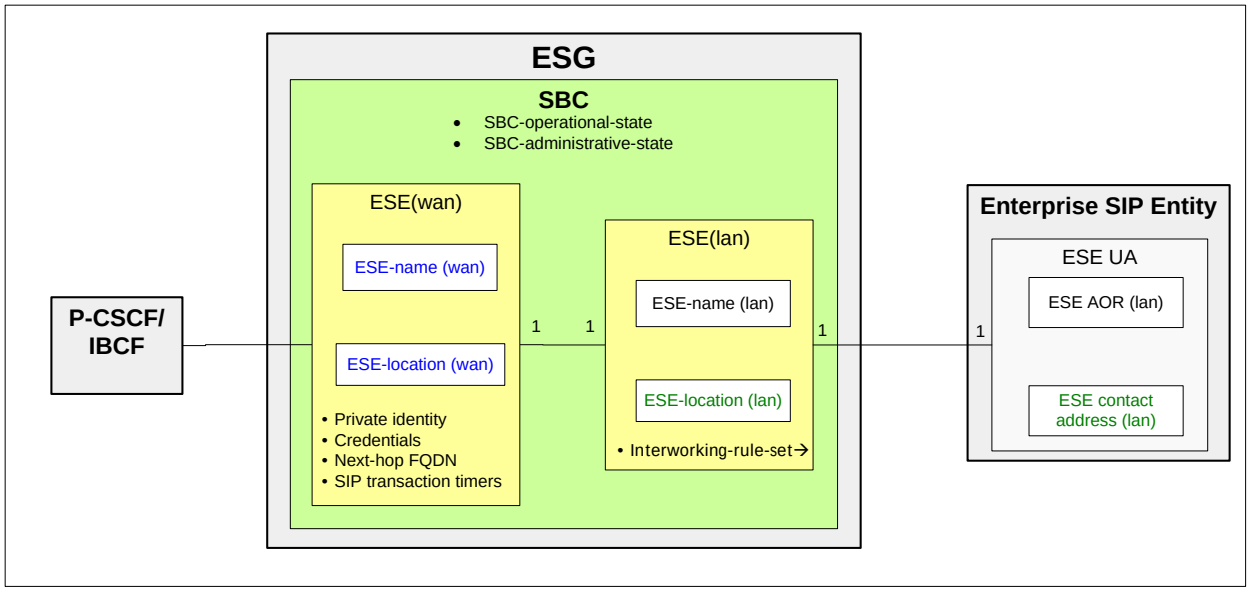

Figure 18 - Extending ESG to B2BUA

A.1.1.3 Adding Support for Multiple Enterprise SIP Entities

Support for multiple ESEs is achieved by creating multiple instances of the ESE(wan)/ESE(lan) objects as shown in Figure 19. These multiple ESE object pairs may reside in a single SBC object instance, or may be spread across multiple SBCs.

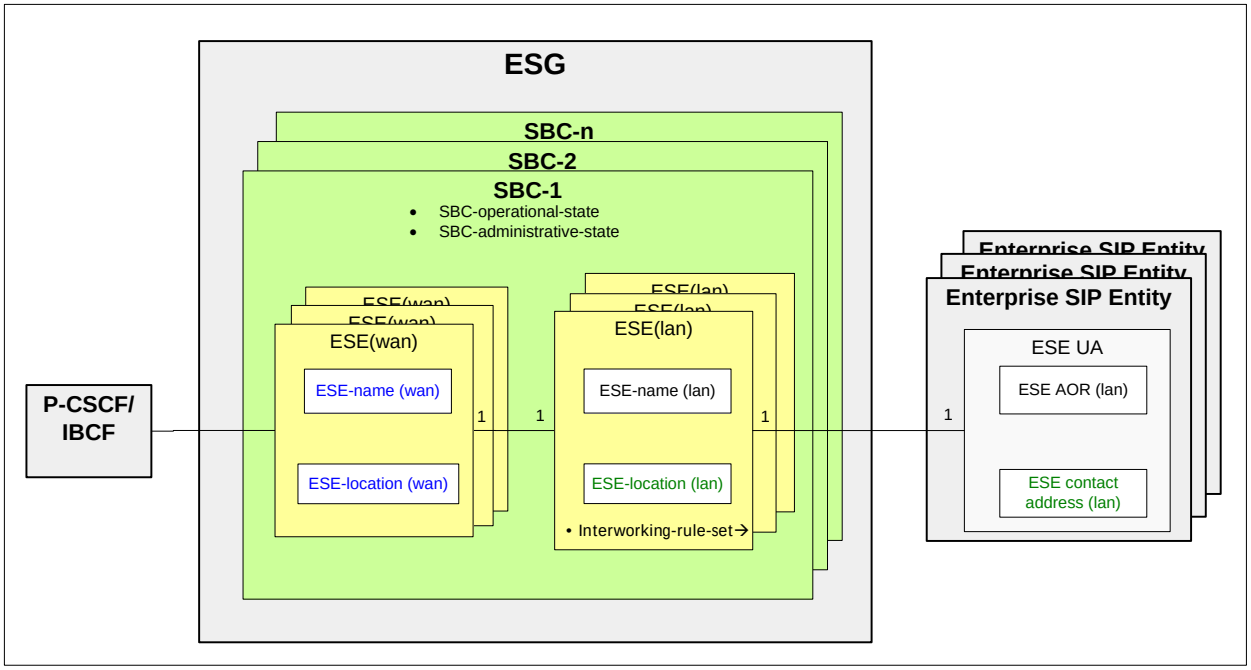

Figure 19 - Support for Multiple ESEs

As shown in Figure 19, each ESE in the Enterprise network is mapped to its own ESE(lan) object in the ESG. This 1:1 mapping is required since each ESE has its own location in the Enterprise network, and therefore there must be a unique ESE(lan) object per ESE to store that location. Figure 19 also shows a 1:1 mapping from ESE(lan) to ESE(wan) object. This form of the data model must be is used when the ESG is serving multiple hosted SIP endpoints, or multiple SIP-PBXs over a SIP Trunk operating in the registration mode. The reason for this is that each registering ESE must have a unique public location on the PC 2.0 network, and therefore each ESE must have its own ESE(wan) to store that location.

When the ESG is serving multiple SIP-PBXs over a static mode SIP Trunk, it can support all SIP-PBXs using a single ESE(wan) object as shown in Figure 20. When this single ESE(wan) instance receives an incoming dialog-initiating SIP request from the PC 2.0 network, it consults a database that maps the E.164 number in the received Request URI to the ESE(lan) object associated with the target SIP-PBX. The E.164 number mapping information can be provisioned locally on the ESG, or the ESG can query an external database such as an ENUM server that returns the ESE-name(lan) for the called E.164 number.

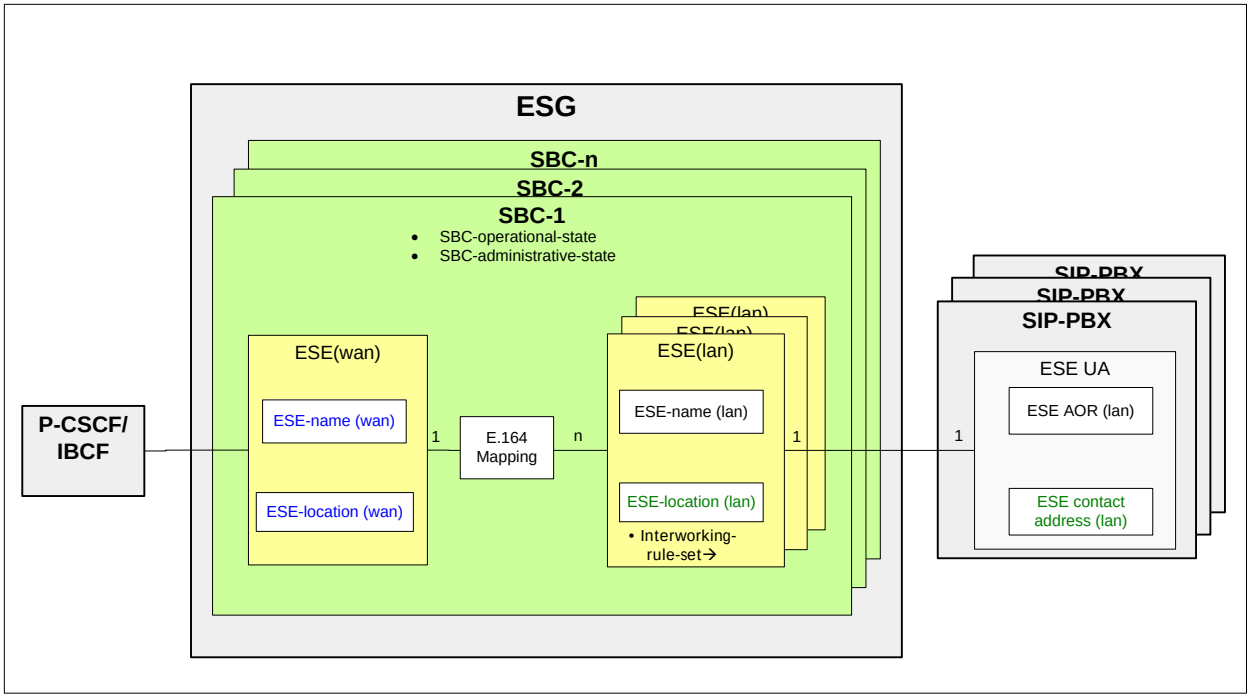

Figure 20 - Support for Multiple ESEs with Single ESE(wan)

A.1.1.4 Adding Support for SIP Proxy

As a configuration option, the SBC may contain a SIP Proxy at the egress/ingress point to/from the PacketCable 2.0 network. This configuration option would most commonly be used when the ESG is serving multiple SIP-PBXs over a static mode SIP Trunk. In this case the Enterprise network appears as a peer network to the Service Provider network, and the proxy serves as an egress/ingress proxy to the "peer" Enterprise network. Although less common, the ESG object model also supports the proxy configuration option when the ESE(wan) registers with the PC 2.0 network. In this case the SBC proxy is inserted in the path and service route as part of the registration procedure.

The SIP Proxy object is contained in the SBC object, and provides routing functions to the ESE(wan) objects contained in that SBC. Figure 21 shows the case where the Proxy object is associated with multiple ESE(wan) objects, where each ESE(wan) object identifies a single SIP-PBX. In the example shown in Figure 21, the PC 2.0 network is configured with a database (e.g., ENUM) that maps the called E.164 number to the target SIP-PBX. The PC 2.0 network places the FQDN of the ESE(wan) associated with the target SIP-PBX in the Request URI of request [1], and the Proxy URI of the target SIP Proxy in the Route header of request [1]. The ESG delivers request [1] to the SIP Proxy identified in the Route header (this ESG could contain additional SIP Proxies associated with other SBC objects). The SIP Proxy removes its URI from the Route header (which exhausts the route set in our example), and then routes the request to the ESE(wan) identified by the FQDN in the Request URI. (Note that in this case the ESE(wan) object does not contain an ESE-location(wan) attribute, since this interface is internal to the ESG.) The [2] INVITE request is routed to the target SIP-PBX via the ESE(lan) linked to the ESE(wan).

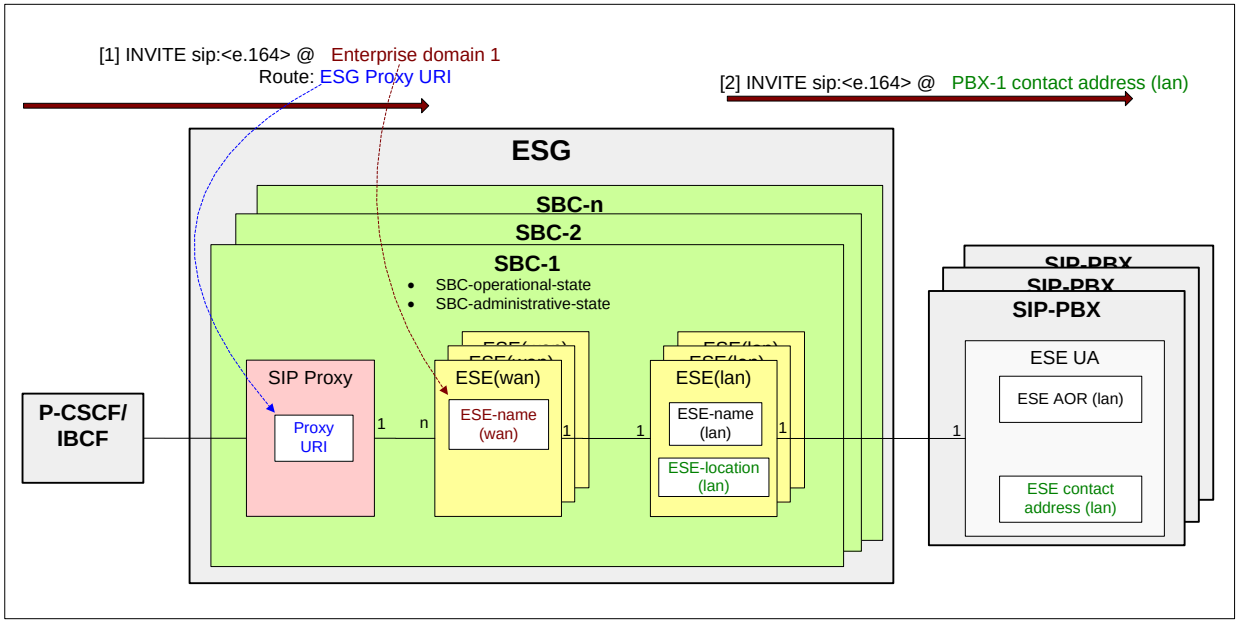

Figure 21 - SBC SIP Proxy Linked to Multiple ESE(wan) Objects

Figure 22 shows the case where the Proxy object is associated with a single ESE(wan) object, that is in turn linked to multiple ESE(lan) objects associated with multiple SIP-PBXs. This configuration option would be used when the PC 2.0 network is configured with a database that is less granular than the previous example; i.e., where the called E.164 number is mapped to the Enterprise domain, and not to a specific SIP-PBX within that domain. The PC 2.0 network places the FQDN of the ESE(wan) associated with the target Enterprise network in the Request URI of request [1], and the Proxy URI of the target SIP Proxy in the Route header of request [1]. The ESG delivers request [1] to the SIP Proxy identified in the Route header. The SIP Proxy removes its URI from the Route header, and then routes the request to the ESE(wan) identified by the FQDN in the Request URI. The ESE(wan) then consults an "E.164 Mapping" database to determine the ESE(lan) associated with the target SIP-PBX, and routes the [2] INVITE request to that PBX.

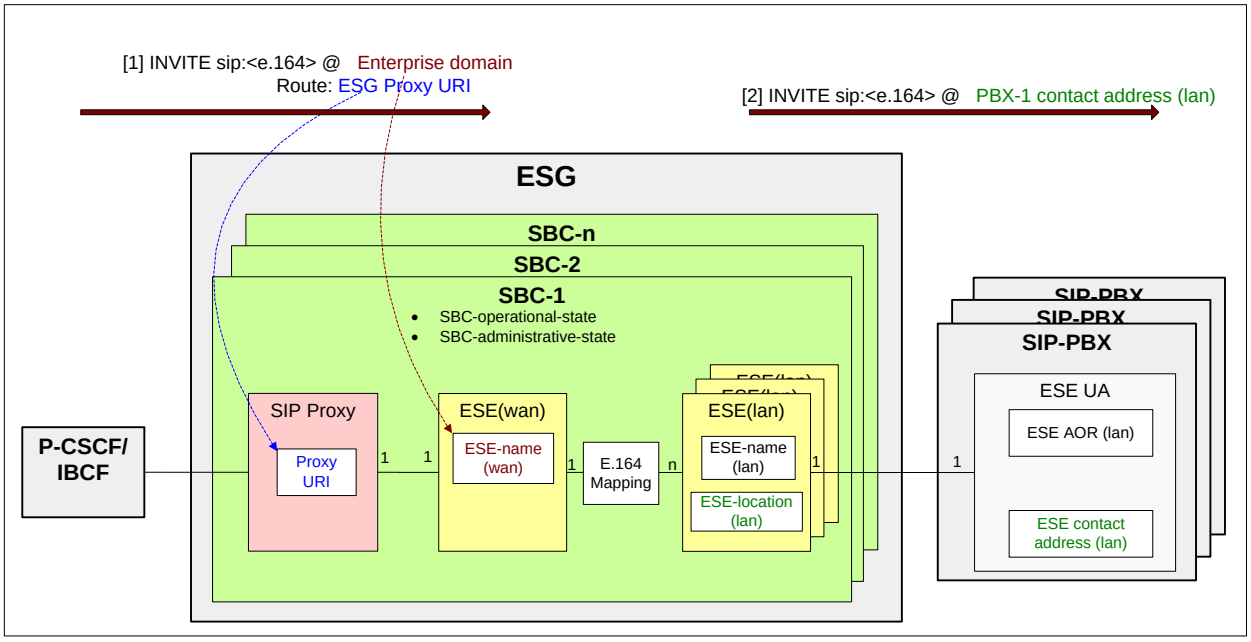

Figure 22 - SBC SIP Proxy Linked to a Single ESE(wan) Object

A.1.1.5 Overview of ESG Request Routing[20]

An enterprise SIP entity is associated with two contact addresses; one LAN address assigned to the SIP entity itself, and one WAN address assigned to the WAN interface of the ESG on behalf of the SIP entity. Therefore, there are two name→location bindings for the SIP entity; one in the LAN IP address space, and one in the WAN IP address space. These location bindings are in turn maintained in two places. The LAN name→location binding is maintained within the ESG, in the ESE(lan) object. The WAN name→location binding is maintained in the Service Provider network (e.g., in the S-CSCF's SIP registrar location database).

The Service Provider network can assign one or more names (Public User identities) to an enterprise SIP entity. For SIP Trunking service, the SIP-PBX is usually assigned a single name (this is referring to the name of the SIP-PBX itself, and not the names of the SIP endpoints serviced by the SIP-PBX). For hosted IP-Centrex service, a SIP endpoint is typically assigned two names; one name containing the public E.164 number of the endpoint, and one name containing the private abbreviated extension number of the endpoint.

Note: Each enterprise SIP entity name can also have two versions; one known to the Service Provider network, and one known to the ESE(lan). The SBC interworks between these two name-versions using the SIP interworking rule set (say, for the case where a SIPconnect1.1-compliant Service Provider network requires a leading '+' character in the user part of the name containing the E.164 number, while the SIP-PBX does not support the leading '+' character). This section assumes that the name assigned in the Service Provider network and the ESE(lan) name are the same (no interworking required).

The Service Provider network is always pre-configured with the enterprise SIP entity names, and it then learns the bound WAN location information for those names from the SIP registration procedure. The ESG can follow this same procedure, where the ESE(lan) is pre-configured with both the public E.164 and private-extension names of the enterprise SIP entity, and it learns the bound LAN location information when the endpoint registers. Alternatively, The ESG ESE(lan) object may not be pre-configured with any information, in which case it has to learn the ESE name(s) and the bound LAN location from the registration procedure. Two cases where the ESE(lan) learns the ESE name(s) are described; one where the ESG is served by an IMS-complaint Service Provider network, and one where the ESG is served by a non-IMS Service Provider network.

1) **IMS-compliant Service Provider network**. Figure 23 shows how an IMS-compliant Service Provider network and ESE(lan) would establish their respective name→location bindings for registering hosted SIP phone (in this example the SIP phone has two names; one containing an E.164 number, and one containing the abbreviated extension number).

[20] Section added by ESG-N-12.0682-8 on 10/28/16 by PO

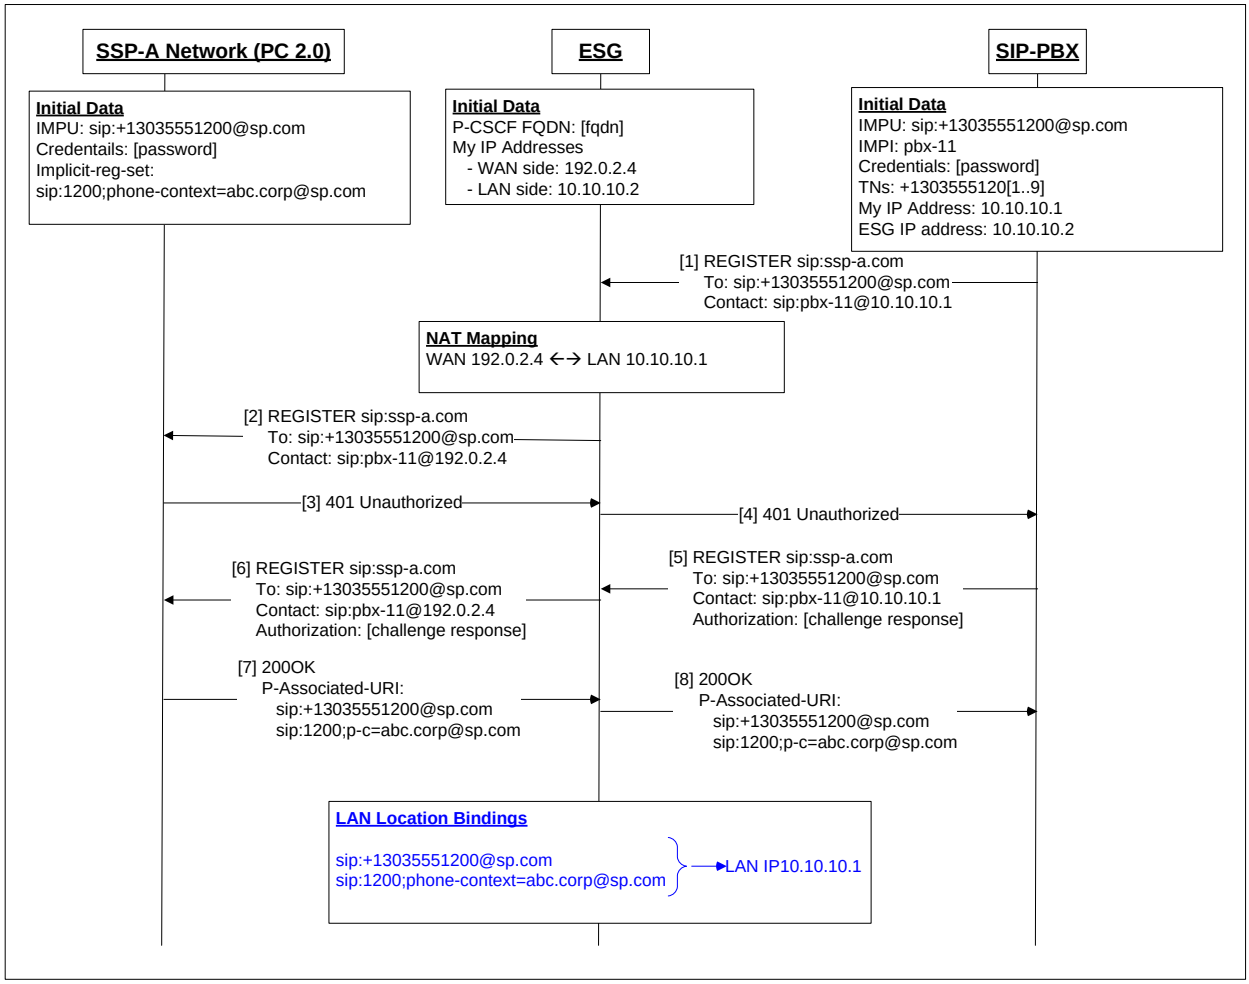

Figure 23 - Discovering LAN Location Bindings

When the ESG receives message [7], it creates two LAN location binding entries for the two Public User Identities contained in the P-Associated-URI header field; one for the name containing the E.164 number, and one for the name containing the abbreviated extension number of the SIP phone. If the P-Associated-URI header contained only a single Public User Identity (say, in the case of a registering SIP-PBX), the ESE(lan) would create a single LAN location binding entry.

2) **Non-IMS Service Provider network:** In this case the Service Provider network does not support the P-Associated-URI header field, and so even though the registering enterprise SIP entity may have multiple names, it returns only the single explicitly registered ESE name in the To: header field of the 200OK response to REGISTER. For the case of a registering SIP phone, where the explicitly registered name contains an E.164 number, and the abbreviated extension number can be derived from the E.164 number (which is usually the case), the ESE can apply a rule configured in the SBC interworking rule set to derive the 2nd private-extension name from the explicitly registered name containing the E.164 number.

Note: This specification does not mandate a specific rule for deriving an associate name from the registering name. The rule could simply consist of a number (e.g., "4" means "the abbreviated number is the right-most 4 digits of the E.164 number"). Or, the rule could be more sophisticated, say to specify a phone-context parameter or other URI parameters associated with the abbreviated number name.

When the ESG is in the normal (non-survivable) mode, the terminating Service Provider network retargets incoming requests to the enterprise SIP entity using the WAN location database; i.e., on receiving a request with a Request-URI containing a Public User Identity of an enterprise user, the terminating Service Provider network follows normal [RFC 3261] routing procedures by retargeting the Request-URI to the WAN location (contact) address bound to the received Public User Identity, before sending the request to the ESG serving the target enterprise SIP entity.

Note: Specific situations may call for alternatives to the basic RFC 3261 retargeting procedure, such as the case of IMS routing to a registered SIP-PBX where either [RFC 6140] or loose-routing is used, or the case where IMS inserts a P-Called-Party-ID header for a hosted UE. But whatever procedure is used, its purpose is to a) deliver the request to the target entity, and b) identify the called user. Here we'll use the term "retarget" to refer to any of these terminating SIP proxy routing procedures.

When the ESG is in survivable mode, it takes on this retargeting function; i.e., when the survivable ESG receives an incoming request (on its LAN interface) with a Request-URI identifying one of its enterprise SIP entities, it uses the ESE(lan) location binding information for the target enterprise SIP entity to retarget the request to the terminating user.

In the case of a SIP-PBX, the survivable ESG needs additional routing information to map the Public User Identity in the incoming Request-URI to the ESE(lan) of the serving SIP-PBX. This information is provided by the E164 Mapping Object defined in Annex A.2.3.7, which maps the E.164 or abbreviated extension number in the user part of the incoming Request-URI to the ESE(lan) of the SIP-PBX.

Note: The ESG uses the registered LAN contact address to retarget requests to registered enterprise SIP endpoints while in survivability. As part of this retargeting process, the ESG may need to perform some special procedure, such as build the Request-URI in a specific manner. For example, if the request is being sent to a SIP-PBX that registered using the "GIN" extension defined in [RFC 6140], and the called user is identified by an E.164 number, then the ESG builds the target URI by setting the user portion to the called E.164 number (including the leading '+' sign), and setting the host portion to the registered LAN contact address. Alternatively, if the target endpoint is an IMS-compatible SIP phone, then the survivable ESG would follow normal RFC 3261 retargeting procedures plus identify the target user in a P-Called-Party-ID header. This specification does not mandate a specific mechanism for determining the retargeting procedure, although one possible option would be to configure the retargeting rules in the interworking rule set of the SBC.

A.1.2 SETA Function Use Cases

The SETA function can be configured to operate in one of two modes; one mode where SETA signaling and media traverse the SBC and therefore verify basic SBC functionality, and an alternate mode where SETA signaling and media are exchanged directly with the PC 2.0 network, bypassing the SBC. The ESG object model to support these two options is illustrated in Figure 24. The SETA object contains the SETA-specific data attributes that control the test call behavior such as test call schedule, destination, and duration. Depending on the type of testing required, the SETA object is linked to different SIP UAs within the ESG to perform the actual SETA functions.

For the case where the SETA line verifies SBC functionality, the SETA object is linked to an ESE(lan) via [1] or ESE(wan) via [2]. The SETA ESE(wan) and ESE(lan) differ from the normal ESE(wan)/ESE(lan) as follows:

- The SETA ESE(lan) does not have an associated Enterprise SIP Entity in the Enterprise network. Instead, the SETA object serves as the Enterprise SIP entity. Any test calls that it originates/terminates are processed according to the Interworking-rule-set configured for ESE(lan).
- The SETA ESE(wan) does not have an associated ESE(lan). Instead, the SETA object serves as the ESE(lan).

For the case where the SETA function bypasses the SBC, the SETA object is linked to a PacketCable 2.0 UE (external from the SBC) via (3). This UE has no analog loop interface to send/receive call signaling and media to/from the user. Instead, the user interface is replaced with the SETA object.

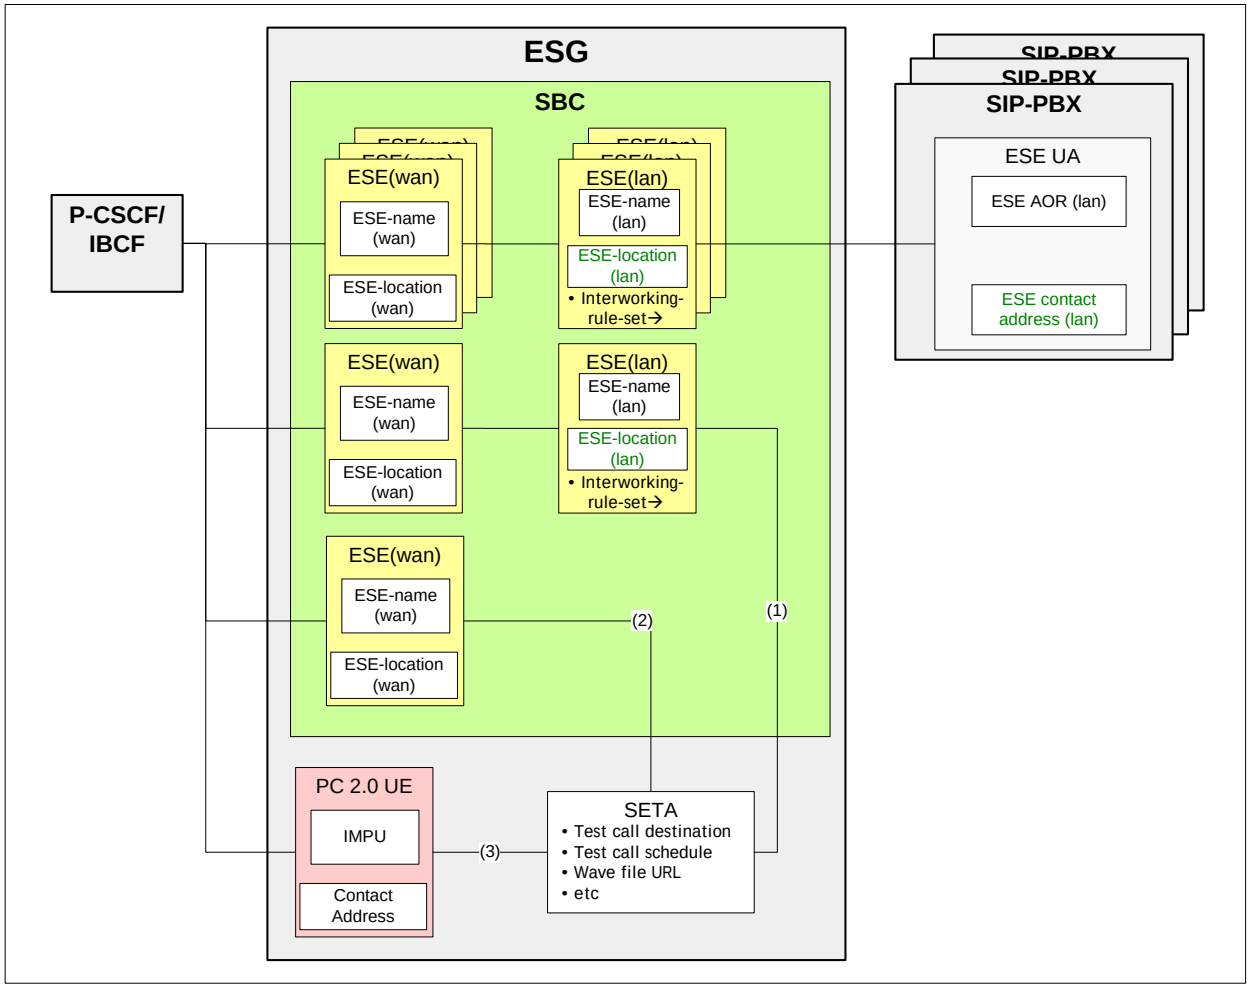

Figure 24 - SETA Line Object Model

A.1.3 Telemetry Function Use Cases

The ESG object model to support the Telemetry function is shown in Figure 25. The Telemetry object contains the Telemetry-specific data attributes that control the collection and reporting of telemetry data. The Telemetry object is associated with the subset of the ESG's ESE(wan) objects for which telemetry data is being collected. The Telemetry object can also be linked to a PacketCable 2.0 UE which can serve as a source for reporting telemetry data to the Service Provider network. The Telemetry object supports configuration options that control whether the PUBLISH of VoIP metrics is reported by an ESE(wan) for which the metrics data was collected, or by the PacketCable 2.0 UE.

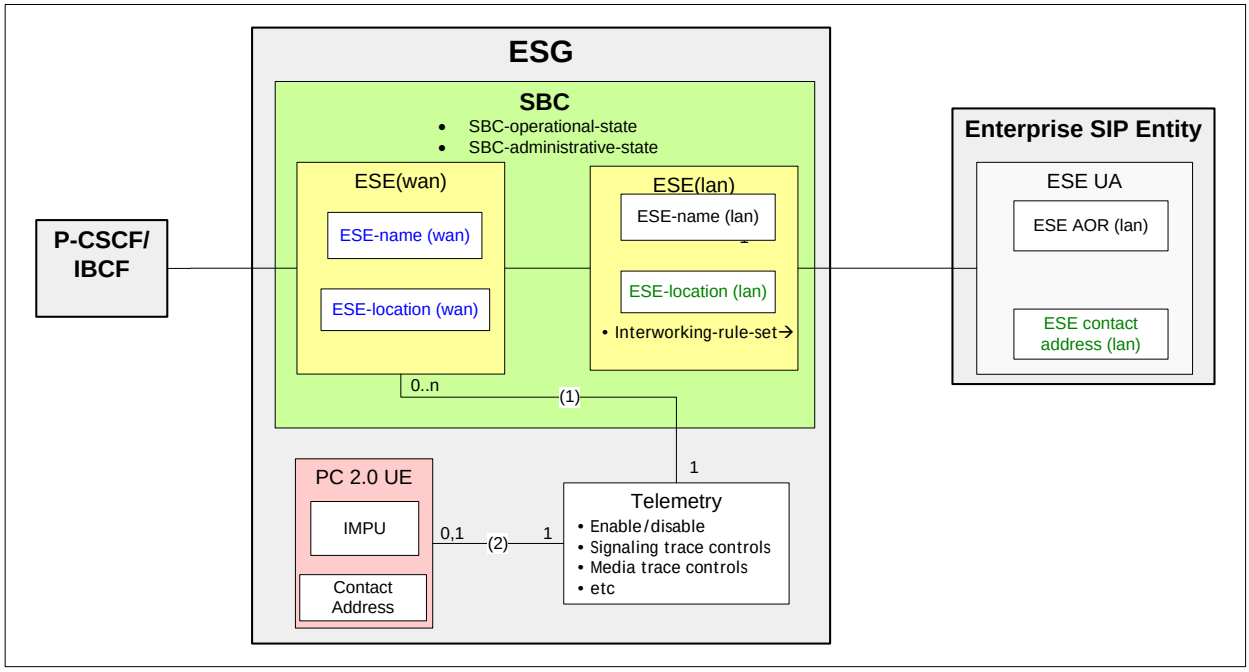

Figure 25 - Telemetry Object Model

The ESG object model defines a global "PUBLISH server" attribute, that defines a separate destination to receive PUBLISH requests containing VoIP metrics. The Service Provider can use this attribute to identify a separate server (separate from the PacketCable 2.0 network) for receiving these VoIP metrics reports.

A.1.4 Provisioning Gateway Function Use Cases[21]

This section describes how the various types of ESG Provisioning Gateways support the basic provisioning message flow for a TR-069-compliant Enterprise SIP endpoint. Specifically, it describes the message flows for CPE-initiated and ACS-initiated CWMP connections via a Traditional NAT, a Twice NAT, and a CWMP ALG.

A.1.4.1 Provisioning Gateway functioning as a NAT

Figure 26 shows a simplified deployment example where the ESG Provisioning Gateway is operating as a NAT (either "Traditional" or "Twice" NAT).

The Enterprise network consists of two subnets; "la" and "lc". The Service Provider network consists of two subnets; "wa" and "wb". If the Provisioning Gateway is operating as a Traditional NAT, then the ESE sets the destination IP address of requests to the ACS or STUN server to the public IP address "wb.1" or "wb.2" of the target server. If the Provisioning Gateway is operating as a Twice NAT, then the ESE sets the destination IP address of requests to the ACS or STUN server to a server-specific port on the ESG's private LAN IP address "lc.1".

[21] New sections added by ESG-N-12.0691-8 on 11/4/16 by PO

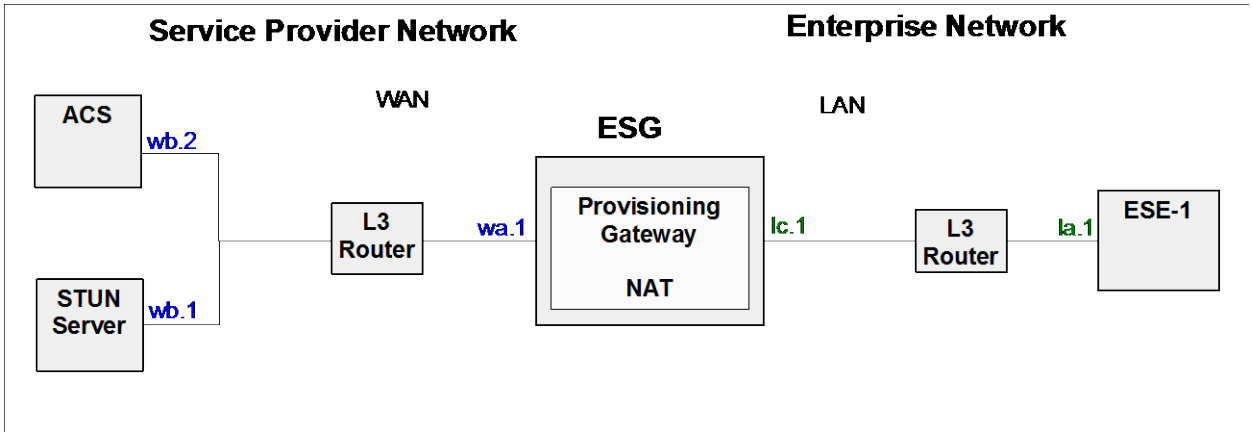

Figure 26 - ESG Data NAT Deployment

A.1.4.1.1 Traditional NAT

This section focuses on how the Traditional NAT routes provisioning traffic between the ESE and the provisioning servers, although the procedures described here would apply to any type of data traffic.

As shown in Figure 27, ESE-1 is configured with the STUN server and ACS URLs which are resolved to the public IP addresses of these servers. The ESE is also configured with the source port for outbound CWMP requests, and the local listening port for its ConnectionRequestURL. The ESG NAT is configured with the LAN IP address:port where it receives incoming requests from the ESE.

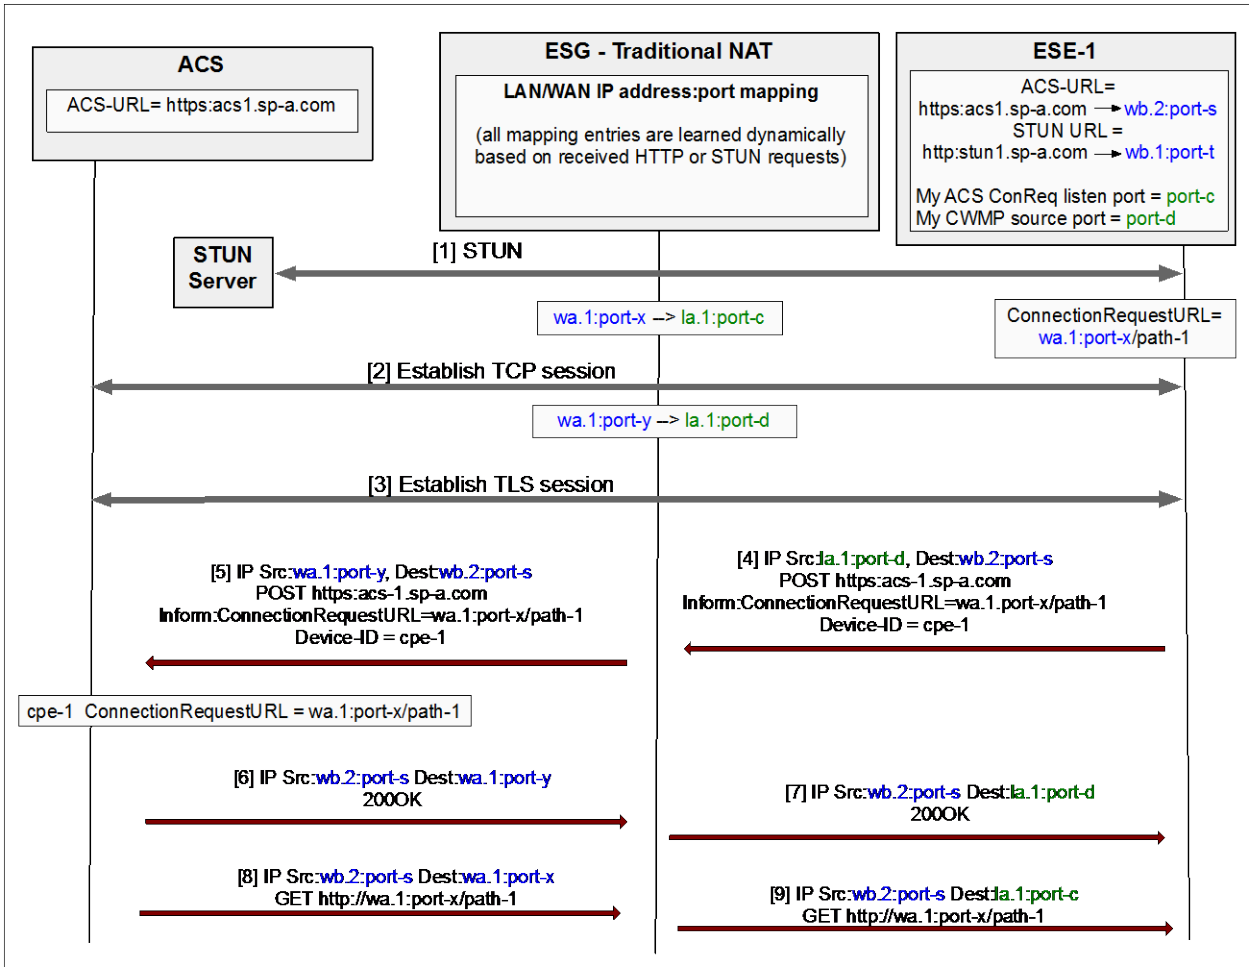

Figure 27 - Traditional NAT Provisioning Flow

[1] The ESE starts sending STUN binding requests from "My ACS ConnReq listen port" to the STUN server via the ESG LAN interface, to discover its WAN ConnectionRequestURL and to open a NAT pinhole for subsequent unsolicited ACS connection requests.

[2] The ESE establishes a TCP connection to the ACS via the ESG LAN interface.

[3] The ESE establishes a TLS session with the ACS via the TCP connection.

[4] The ESE sends a connection request Inform message via the TCP/TLS connection. On receiving [4], the ESG NAT updates the source IP address:port per normal/traditional NAT procedures.

[5] The ACS saves the received ConnectionRequestURL for ESE device cpe-1.

[6] The ACS responds with a 200OK.

[7] The ESG NAT updates the response destination IP addresses and forwards the response to the ESE.

[8], [9] Later, the ACS sends a UDP connection request to ESE-1, to the previously reported ConnectionRequestURL. The ESG NAT relays the request to ESE-1 based on the NAT bindings opened by the periodic STUN requests.

A.1.4.1.2 Twice NAT

This section focuses on how the Twice NAT routes provisioning traffic between the ESE and the provisioning servers.

Figure 28 shows the initial CWMP connection procedure initiated by ESE-1, and the resulting unsolicited ACS connection request from the ACS. ESE-1 is statically configured to resolve the STUN server and ACS URLs to an IP address:port on the ESG LAN interface. The Twice NAT is configured with its LAN listening IP address:ports for incoming STUN and CWMP requests from the ESE. The STUN listening port is mapped to the actual STUN server's public IP address:port, and the CWMP listing port is mapped to the ACS's public IP address:port. The Twice NAT creates address mapping entries dynamically based on received HTTP and STUN binding requests, the same as a Traditional NAT.

If the ACS can redirect the ESE to other servers located within the Service Provider network, then the ESE and Twice NAT configuration data entries shown here to resolve and map the ACS address would have to be repeated for these other servers.

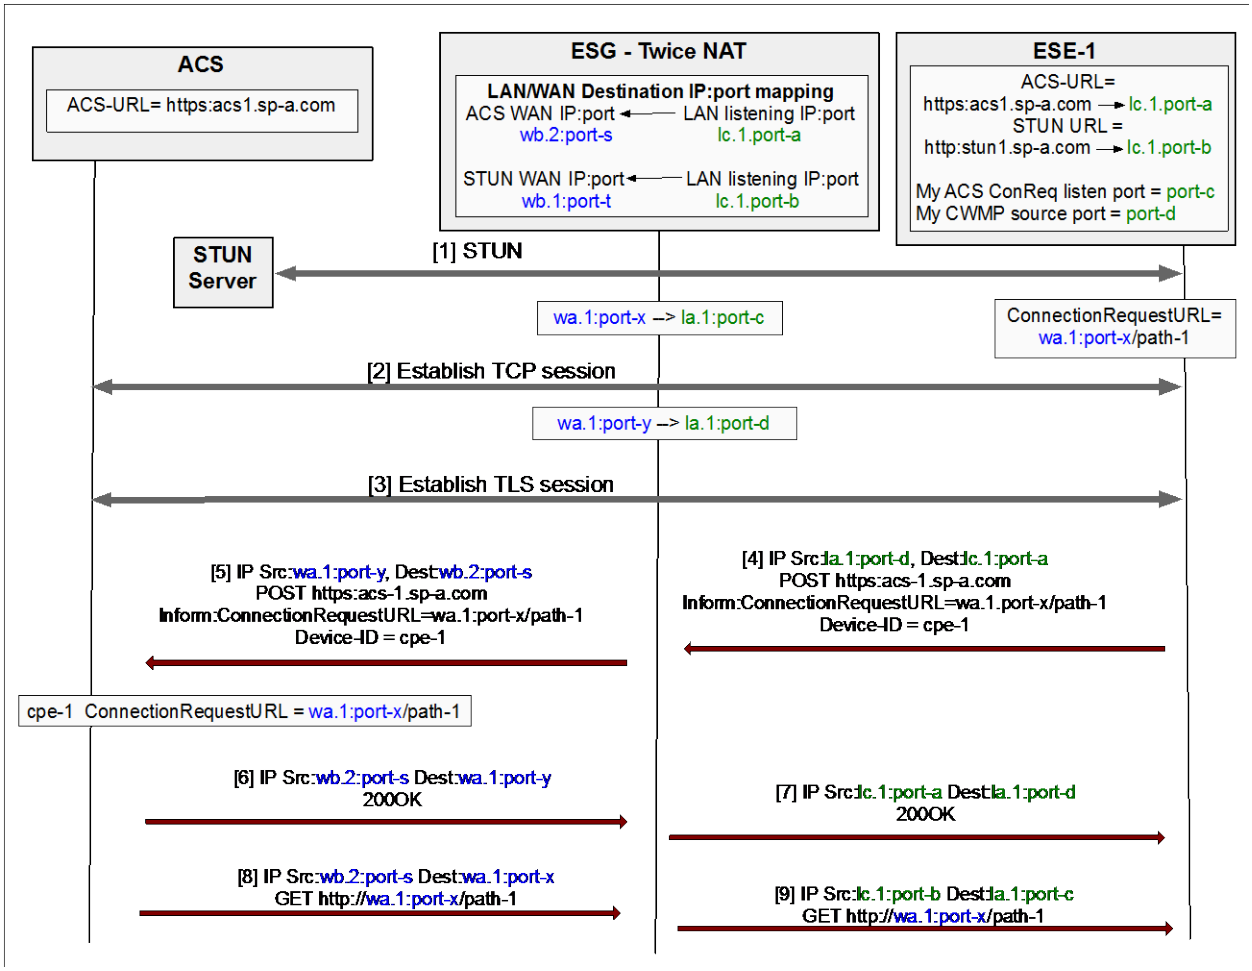

Figure 28 - Twice NAT Provisioning Flow

[1] The ESE starts sending STUN binding requests via ESG LAN address lc.1:port-b to discover its WAN ConnectionReqURL and to open a NAT pinhole for subsequent unsolicited ACS connection requests. On

receiving the STUN binding request from the ESE on port-b of its LAN interface, the ESG Twice NAT retargets the request to the public IP address of the STUN server.

[2] The ESE establishes a TCP connection to the ACS via the ESG LAN address 1c.1:port-a. On receiving the TCP connection-establishment request on port-a of its LAN interface, the ESG Twice NAT retargets the request to the locally configured public IP address of the ACS.

[3] The ESE establishes a TLS session with the ACS via the TCP connection. On receiving the TLS session-establishment requests on port-a of its LAN interface, the ESG Twice NAT retargets the request to the locally configured public IP address of the ACS.

[4] The ESE sends a connection request Inform message via the TCP/TLS connection. On receiving [4], the ESG Twice NAT updates the source IP address:port per normal/traditional NAT procedures, and updates the destination IP address:port based on the configured mapping table.

[5] The ACS saves the received ConnectionRequestURL for ESE-1.

[6] The ACS responds with a 200OK.

[7] The ESG NAT updates the response source/destination IP addresses and forwards to the ESE.

[8], [9] Later, the ACS sends a UDP connection request to ESE-1, based on the saved ConnectionRequestURL. The ESG NAT relays the request to ESE-1 based on the NAT

A.1.4.2 Provisioning Gateway functioning as a CWMP ALG

Figure 29 shows the CWMP ALG within the overall provisioning architecture defined in [TR-069]. The Enterprise SIP Entity, acting as a TR-069 CPE, obtains its configuration data from and is managed by the ACS located in the Service Provider network. The CWMP ALG sits at the boundary between the Enterprise and Service Provider networks to perform application-level WAN-LAN interworking for the CWMP protocol running between the ESE and its ACS.

As shown in the diagram, an ESG can support multiple CWMP ALG instances, where each ALG instance provides interworking between a group of ESE devices and the single ACS serving those devices.

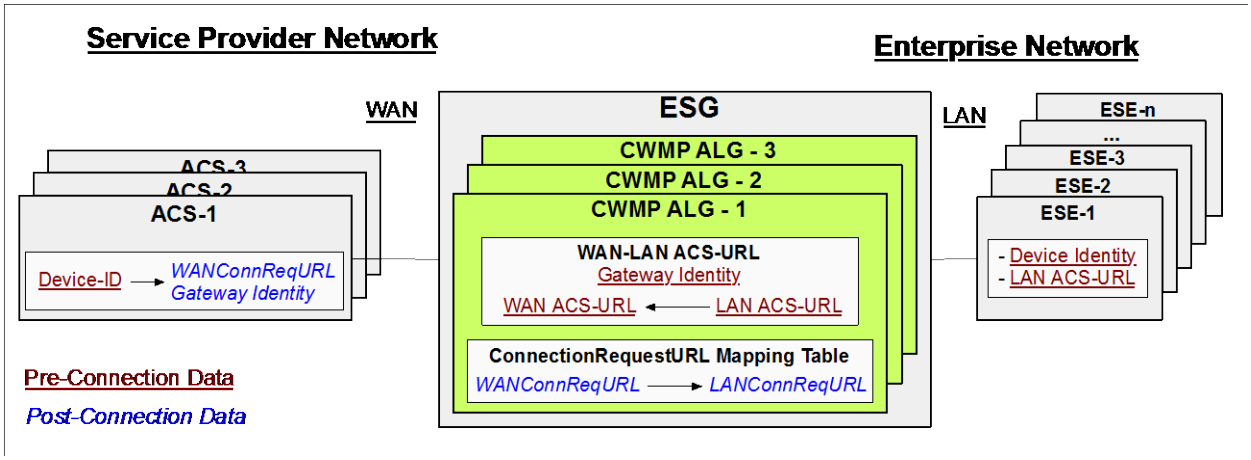

Figure 29 - CWMP ALG within the TR-069 Architecture

A.1.4.2.1 CWMP ALG Data

Figure 29 identifies the configured data items (red-underscore font) known by the CWMP ALG and its associated TR-069 entities before the ESE first establishes a connection with its ACS. It also shows the dynamic data (*blue-*

italics font) that is learned during CWMP connection establishment and maintained by the ACS and the CWMP ALG once the ESE has established a connection.

Pre-Connection Data

Figure 29 highlights the data items known to the various entities before the ESE first establishes a connection with its ACS. These data items are either statically configured, or learned dynamically (say via DHCP).

The ESE initially knows its own device identity, and the URL of its ACS. In this case the ACS URL points to the ESG LAN interface.

The CWMP ALG knows its Gateway Identity (which is the ESG's Device Identity). It also knows the LAN-side ACS URL, and uses this URL to identify the listening port for CWMP connection requests from the ESEs. The CWMP-ALG also knows the publically routable WAN-side ACS URL. The CWMP ALG is not initially configured with the identity of the ESE devices that it serves. Instead, it learns the Device Identity of each ESE when the ESE first establishes a connection with its ACS.

The ACS initially knows the identity of the devices it serves.

Post-Connection Data

Figure 29 also shows the transient data that is learned and maintained by the ACS and the CWMP ALG when the ESE first establishes a connection with the ACS.

From the ESE's perspective, the ESG is its ACS. Since there is no NAT between the ESE and the ESG, the ESE reasonably assumes that the ConnectionRequestURL it populates in the Inform request sent to the ESG is globally routable. The CWMP ALG knows better - it understands the WAN/LAN network topology, and knows that ConnectionRequestURL from the ESE is routable only within the private Enterprise network. Therefore, the CWMP ALG assigns a new globally routable WAN ConnectionRequestURL that resolves to the ESG's WAN IP address. The CWMP ALG replaces the LAN-based ConnectionRequestURL received in the Inform request from the ESE with the new WAN ConnectionRequestURL, and adds its own Gateway Identity to the Inform, before forwarding the Inform to the ACS. The CWMP ALG also maintains the mapping between the WAN and LAN ConnectionRequestURL, so it can route subsequent connection requests from the ACS to the correct ESE.

Note: In addition to the data shown in Figure 29, for each ConnectionRequestURL mapping entry the CWMP ALG also saves the device identity information contained in the Inform request received from the ESE. The CWMP then uses the saved device identity to detect when a WAN/LAN ConnectionRequestURL mapping entry becomes stale - say when a new device reuses an already-mapped ConnectionRequestURL.

In receiving the CWMP Inform request to establish a connection, the ACS saves the received WAN ConnectionRequestURL, which it can then use to send subsequent connection requests via the ESG to the ESE. It also saves the Gateway Identity associated with the ESE Device Identity.

A.1.4.2.2 CWMP ALG Procedures

Figure 30 shows a Hosted IP Centrex deployment example where the ESG LAN interface and the SIP phones are located across multiple subnets with the Enterprise network. ESE-1 and ESE-2 are configured with an ACS address pointing to the LAN interface of the ESG. ESE-1, ESE-2, and the ESG LAN interface are assigned IP addresses la.1, lb.1, and lc.1, respectively. The ESG WAN interface is assigned IP address wa.1. ACS-1 contains a Session Receiver at IP address wb.2 that receives CWMP Inform requests from the ESEs, and a Connection Requester that initiates unsolicited CWMP connection requests to the ESEs. The CWMP ALG is configured with the actual public URL of the ACS ("acs-1", in this example).

The actual TR-069 provisioning procedures are performed within the context of a "connection" initiated by the CPE. For example, a newly rebooted CPE establishes a connection with the ACS (by sending a CWMP Inform request) so

it can be downloaded by the ACS. Since CWMP is carried in HTTP messages, the initial Inform request from the CPE (carried in an HTTP POST request) opens a pinhole through the NAT to enable the subsequent CWMP conversation between the CPE and ACS to take place. The conversation happens over this connection established by the Inform request. Once the conversation is done, the connection is released.

Sometimes, the ACS wants to initiate a conversation with the CPE when there is no connection. It does this by sending a special unsolicited "connection request" to the CPE, which essentially prompts the CPE to establish a new connection with the ACS. Since the ACS "connection request" is contained in an unsolicited HTTP request, it is blocked by most firewalls, and is unable to traverse NATs. The CWMP-aware ALG has application-level knowledge, and therefore knows to route this unsolicited connection request message from the ACS to the target CPE, based on WAN-to-LAN address mapping information that was established during a previous connection initiation from the CPE.

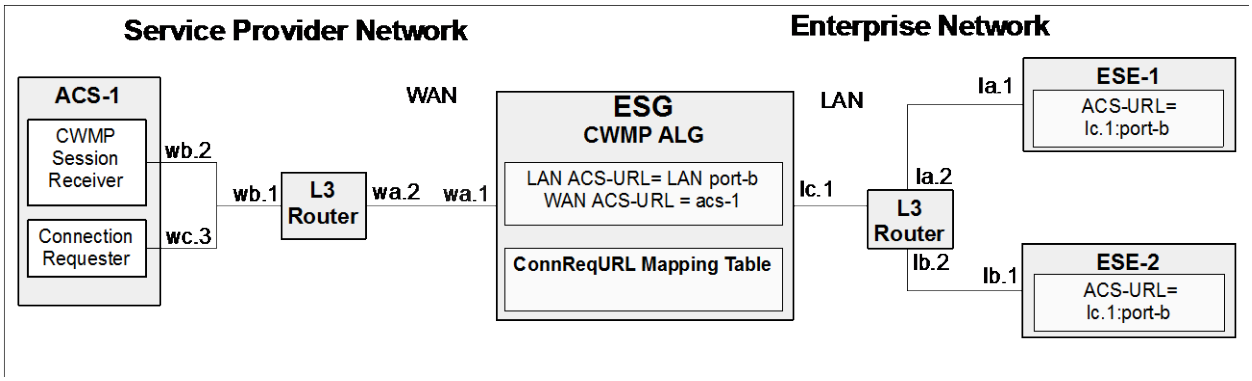

Figure 30 - CWMP Deployment Example

Figure 31 shows the CWMP message flow for an ESE-initiated connection request, and a subsequent ACS-initiated unsolicited connection request for the deployment example shown in Figure 30.

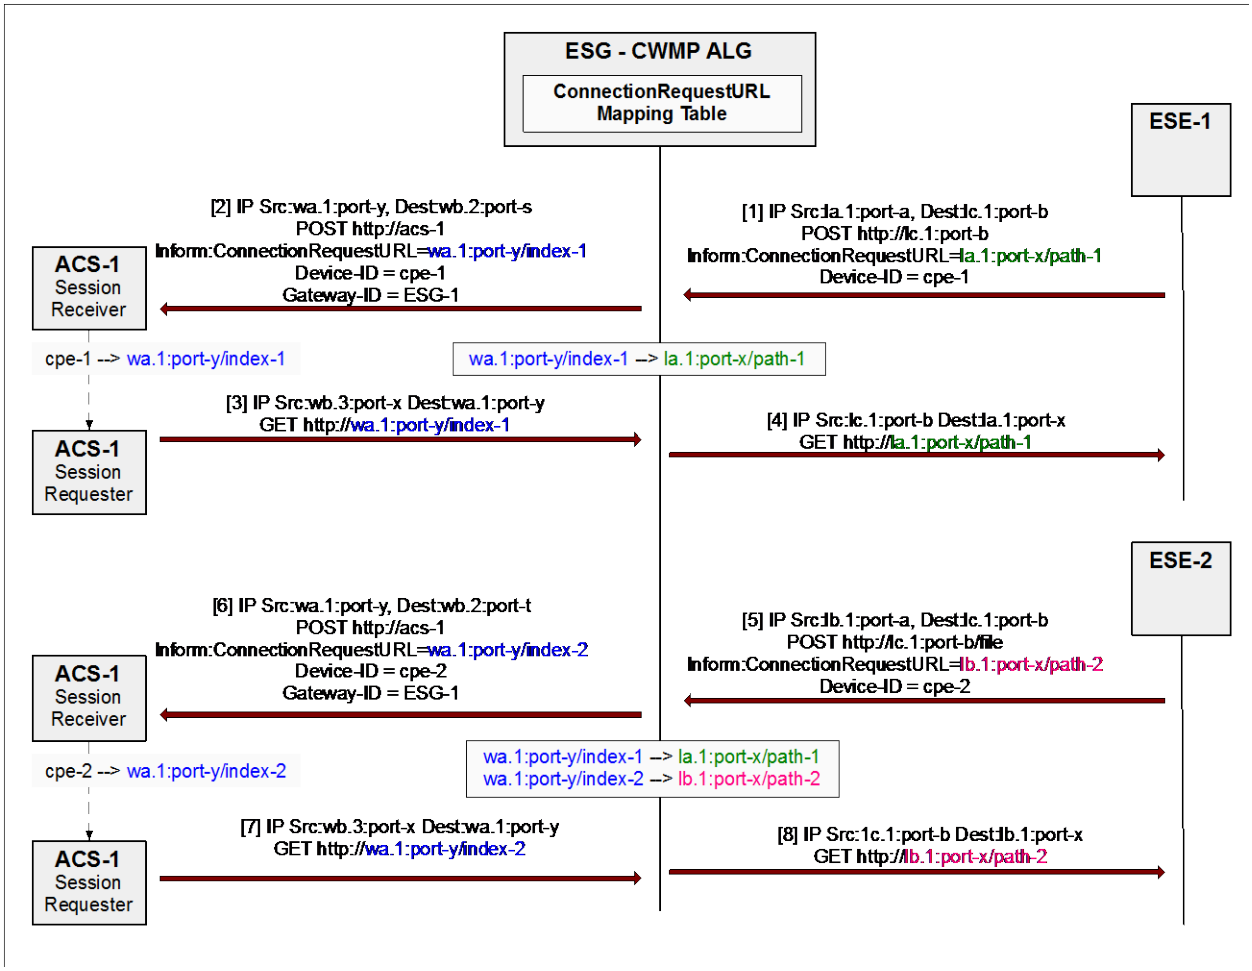

Figure 31 - CWMP Message Flow Example

[1] At reboot time ESE-1 sends an Inform request to its configured ACS (which in this case is the LAN interface of the ALG). The Inform contains the ESE's device identity, and a ConectionRequestURL identifying the ESE's LAN address:port that is listening for subsequent unsolicited connection requests from the ACS.

[2] On receiving [1], the CWMP ALG assigns a WAN IP address:port/path mapped to the ESE's LAN address:port/path, and updates the ConectionRequestURL with the newly assigned WAN IP address:port/path, and forwards the request on to ACS-1. The ALG stores the WAN-to-LAN address mapping in order to route subsequent requests from the ACS to ESE-1. The ALG also stores the Device-ID, which helps detect the case where the current WAN-to-LAN address mapping becomes stale. On receiving [2], the ACS stores the connection URL for the identified device.

[3] Later (e.g., hours or days later), ACS-1 Connection Requester initiates a connection request to ESE-1 by sending an HTTP GET to the ConnectionRequestURL (i.e., to the ALG WAN address:port that is listening for ACS requests). On receiving [3], the ALG interworks the IP header fields and request URI values from WAN-to-LAN based on the previously saved address mapping info.

[4] The CWMP ALG forwards the request to the target ESE-1. At this point ESE-1 initiates a connection with the ACS to perform whatever action is requested by the ACS.

[5], [6], [7], [8], repeats the process for a 2nd ESE - ESE-2. This establishes a 2nd WAN-to-LAN address:port mapping in the ESG for mapping subsequent requests from the ACS to ESE-2.

A.2 ESG Object Model Definitions

A.2.1 ESG Object Model Data Types

There are no data types defined in this module.

A.2.2 ESG Object Model Class Diagram

The ESG object model class diagram shown in Figure 32 covers the following three areas:

- SBC Configuration
- SETA Configuration
- Statistics and Debugging Information

The SBC Configuration defines the relationships and data attributes of the objects contained in the SBC. Basically, an SBC is a container for multiple Enterprise SIP Endpoints (ESEs). As shown in Figure 32, the "SBCCfg" object sits at the top of the hierarchy. It points to one or more "WanESE" objects which contain the public address information of the Enterprise SIP Entities served by that SBC. If the "WanESE" is acting as a SIP UA, then it points to the UE data model defined in PacketCable 2.0. Each "WanESE" object also points to one or more "LanESE" objects which represent the individual Enterprise SIP Entities within the Enterprise network that are served by the containing SBC. Each "LanESE" identifies the interworking rule-set and the firewall rule-set required by its associated Enterprise SIP Entity. Finally, the "SBCCfg" object is linked to the "SBC Mapping Status" object that contains multiple read-only attributes describing each "LanESE"→"WanESE" pair.

The SETA Configuration defines the objects and their attributes to support the SETA function.

The Statistics and Debugging Information defines the collection of objects and attributes required to support the Telemetry function. It identifies the objects used to control and collect the call completion statistics, crossing-threshold events on calls quality and reliability, VoIP metrics, and call signaling and RTP payload trace files, and all other debug information.

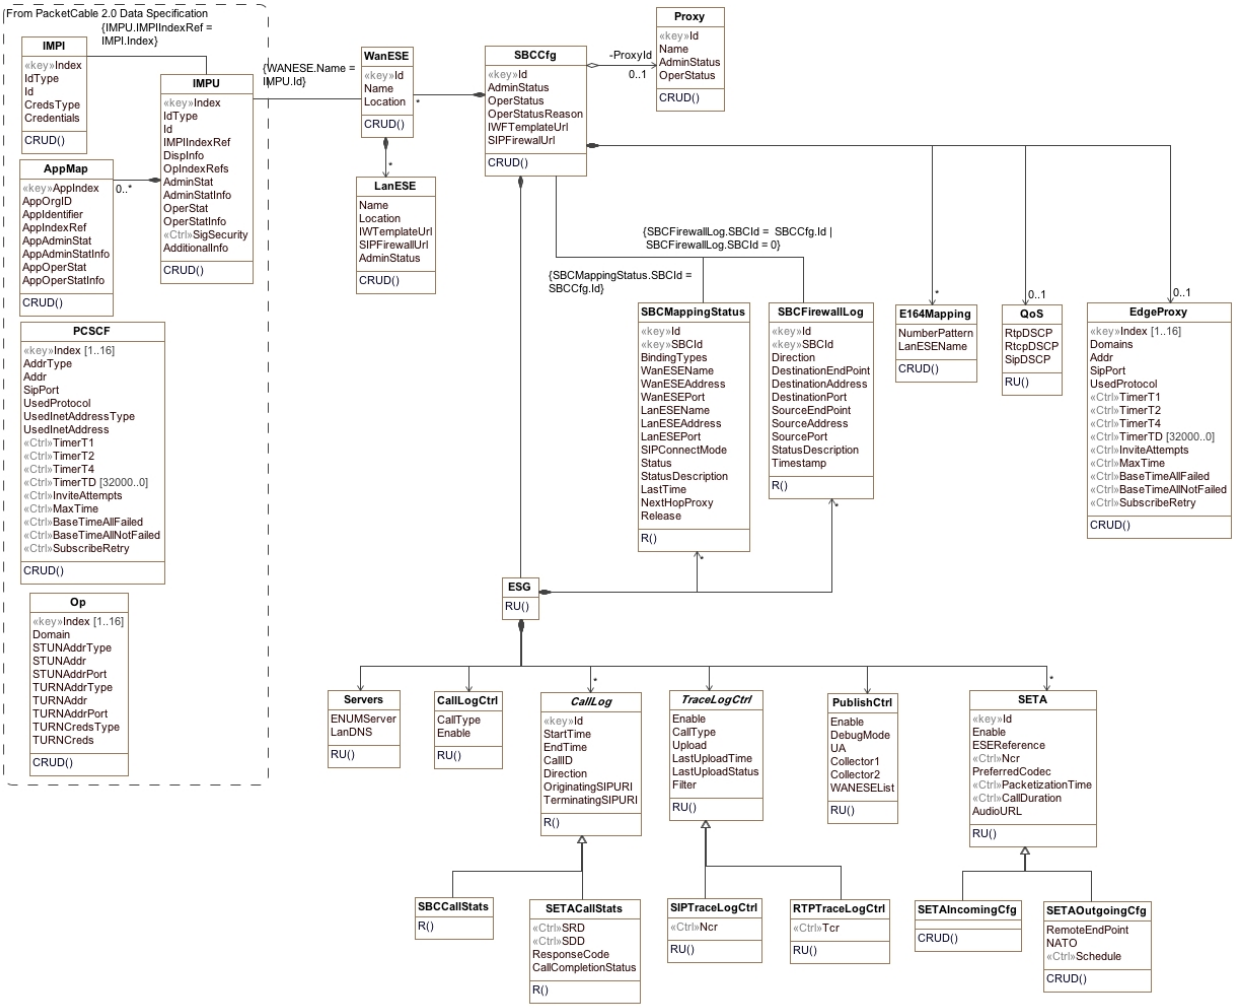

Figure 32 - ESG Object Model Diagram[22]

A.2.3 ESG Object Model Description

A.2.3.1 SBCCfg Object[23]

This object represents the Session Border Controller configuration of the ESG. An incoming request is associated with a SBC by first matching the calling party (for requests incoming from the Enterprise network) or called party (for requests incoming from the SP network) to the SBC ESEs. SBCs are matched in ascending order of the Id key.

Object Operations:

None

[22] ESG-N-12.0691-8 10/28/16 by PO
[23] ESG-N-12.0691-8 10/28/16 by PO

Table 3 - SBCCfg Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	Key			
AdminStatus	Enum	CRUD	inService(1),outOfServcice(2),outOfServiceIdle(3)		
OperStatus	Enum	R	inService(1),outOfServcice(2) survivable(3)		
OperStatusReason	AdminString	R			
IWFTemplateUrl	AdminString	RU			
SIPFirewallUrl	AdminString	RU			

Attribute Descriptions:

- Id

  This key represents the unique identifier of an object instance.

- AdminStatus

  This attribute represents the administrative state of the SBC set by the operator. It has the following values:

  'inService' indicates the SBC is administratively active.

  'outOfService' indicates the SBC is administratively inactive.

  'outOfServiceIdle' indicates the SBC will transition to 'outOfService' once the number of active calls drops to zero. New call attempts are blocked in this state. On transitioning to this state, the ReturnToService timer is started. If the timer expires before the number of active calls drops to zero, then the SBC is set back to 'inService'

- OperStatus

  This attribute represents the operational state of the SBC (i.e., indicates whether or not the SBC is working). It has the following values:

  'inService' indicates that the SBC is operational.

  'outOfService' indicates that the SBC is not operational.

  'survivable' indicates that the SBC is in survivable mode; i.e., the SBC has lost connectivity with the Service Provider network, but is still operational and can support intra-enterprise calls

- OperStatusReason

  This attribute represents a human readable reason, or explanation of the current operating status of the SBC.

- IWTemplateUrl

  This attribute defines the location of a file for the Interworking function applied to all LanESEs, unless a particular LanESE has its own template file.

- SIPFirewallUrl

This attribute defines the location of a file for the SIP firewall applied to all LanESEs, unless a particular LanESE has its own template file.

A.2.3.2 WanESE Object

This object represents the WAN ESE. The associations of WanESE to LanESE(s) serve as a SIP-Aware Access Control List (ACL);i.e., the SBC discards messages not related to these relationships.

Object Operations:

None

Table 4 - WanESE Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
Name	AdminString	CRUD	SIZE(0..256)		
Location	AdminString	CRUD	SIZE(0..256)		

Attribute Descriptions:

- Id

This key represents the unique identifier of an object instance.

- Name

This attribute represents the publicly known identity of the ESE. It may identify the AOR of a SIP endpoint or a SIP-PBX operating in the registration mode, or an FQDN of a SIP-PBX operating in the static mode.

- Location

This attribute defines the host IP address and port associated with the ESE. The format of this attribute is a variation of the ABNF notation of the hostport part of SIP URI from [RFC 3261]:

hostport = host [":" port]

host = hostname / IPv4address / IPv6reference

A hostname may be used when DNS resolution is available, otherwise IP Address notation is preferred.

A.2.3.3 LanESE Object

This object represents the LAN ESE.

Object Operations:

None

Table 5 - LanESE Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Name	AdminString	CRUD	SIZE(0..256)		
Location	AdminString	CRUD	SIZE(0..256)		
IWTemplateUrl	Uri	CRUD	SIZE(0..256)		
SIPFirewallUrl	Uri	CRUD	SIZE(0..256)		
AdminStatus	Enum	CRUD	active(1),inactive(2)		

Attribute Descriptions:

- Name

This attribute represents the AOR of a SIP-PBX or SIP endpoint. A comma separated list is used to define more than one identity (e.g., a SIP URI defining an E.164 number, and a SIP URI defining the abbreviate dial number of the end point).

- Location

This attribute defines the host IP address and port associated with the ESE. The format of this attribute is a variation of the ABNF notation of the hostport part of SIP URI from [RFC 3261]:

hostport = host [":" port]

host = hostname / IPv4address / IPv6reference

A hostname may be used when DNS resolution is available, otherwise IP Address notation is preferred.

- IWTemplateUrl

This attribute defines the location of a file for the Interworking function.

- SIPFirewallUrl

This attribute defines the location of a file for the SIP firewall.

- AdminStatus

This attribute represents the administrative state of the LAN ESE.

A.2.3.4 Proxy Object

This object represents a SIP proxy associated with a SBC.

Object Operations:

None

Table 6 - Proxy Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
Name	AdminString	CRUD	SIZE(0..256)		

Attribute Name	Type	Access	Type Constraints	Units	Default
AdminStatus	Enum	CRUD	inService(1),outOfService(2), outOfServiceIdle(3)		
OperStatus	Enum	CRUD	inService(1),outOfService(2)		

Attribute Descriptions:

- Id

This key represents the unique identifier of an object instance.

- Name

This attribute identifies the SIP URI of this SIP Proxy. This is used to route SIP request to this object (i.e., based on the SIP URI contained in the Route header field of a received SIP request).

- AdminStatus

This attribute represents the administrative state of the SIP Proxy.

- OperStatus

This attribute represents the operational state of the SIP Proxy.

A.2.3.5 EdgeProxy Object

This object represents the information of the next hop used in the signaling path by an endpoint. This information is used by ESEs (UAs) configured to operate in the static mode defined in [SIPconnect1.1].

Object Operations:

None

Table 7 - EdgeProxy Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Index	unsignedInt	key			
Domains	AdminString	CRUD			
Addr	InetAddress	CRUD			""
SipPort	InetPortNumber	CRUD			
UsedProtocol	PktcEUEDevSipProtID	R			
TimerT1	unsignedInt	CRUD		milliseconds	500
TimerT2	unsignedInt	CRUD		milliseconds	4000
TimerT4	unsignedInt	CRUD		milliseconds	5000
TimerTD	unsignedInt	R		milliseconds	32000
InviteAttempts	unsignedInt	CRUD	1..7	attempts	

Attribute Name	Type	Access	Type Constraints	Units	Default
MaxTime	unsignedInt	CRUD		seconds	
BaseTimeAllFailed	unsignedInt	CRUD		seconds	
BaseTimeAllNotFailed	unsignedInt	CRUD		seconds	
SubscribeRetry	unsignedInt	CRUD		seconds	

Attribute Descriptions:

- Index

This key represents the unique identifier of an object instance.

- Domains

This attribute represents a comma-separated list of domains (sub-domains and FQDNs) using this proxy information. A SIP Entity extracts its domain from its own name, and matches the first instance in the EdgeProxy object that contains such domain. the SIP Entity may find additional instances matching same domain and uses them as alternate next hops.

- Addr

This attribute represents the address of the proxy.

- SipPort

This attribute contains a SIP Port the edge proxy is listening. By default port 5060 is defined for SIP udp/tcp transports and 5061 for tls.

- UsedProtocol

This attribute contains the SIP Protocol used.

- TimerT1

This attribute represents the SIP Timer T1, an estimate for the round trip time in the system. Please refer to [PKT 24.229] for more information.

- TimerT2

This attribute represents the SIP Timer T2, an estimate for the maximum retransmit interval for non-INVITE requests and INVITE responses. Please refer to [PKT 24.229] for more information.

- TimerT4

This attribute represents the SIP Timer TD, indicates the wait time for response retransmits. Please refer [PKT 24.229] for more information.

- TimerTD

This attribute represents the SIP Timer TD, an estimate for the maximum duration a message will remain in the network. Please refer to [PKT 24.229] for more information. If the protocol used for a SIP Session is UDP this value is used for SIP Timer D, otherwise is ignored and the SIP session.

- InviteAttempts

This attribute represents the total number of INVITE message attempts before the SIP transaction is considered as failed due to no response.

The total Timer TB MUST be derived from the total number of SIP INVITE message attempts as follows:

TB = [2^(n -1) - 1]*T1

n: total number of INVITE attempts

T1 = Timer T1

For example, if the number of INVITE attempts is 3, (initial INVITE + 2 retries)

TB = [(1 - 1)^2 + (2 - 1)^2 + (3 - 1)^2]*0.5 = 3.5 secs.

Please refer to [PKT 24.229] for more information.

- MaxTime

  This attribute represents the 'max-time' SIP Registration Recovery Timer as defined in [RFC 5626]. Please refer to [PKT 24.229], and [RFC 5626] for more information.

- BaseTimeAllFailed

  This attribute represents the 'base-time (if all failed)' SIP Registration Recovery Timer as defined in [RFC 5626]. Please refer to [PKT 24.229], and [RFC 5626] for more information. If the protocol used for a SIP Session is UDP this value is used for SIP Timer D, otherwise is ignored and the SIP session.

- BaseTimeAllNotFailed

  This attribute represents the 'base-time (if all have not failed)' SIP Registration Recovery Timer as defined in [RFC 5626]. Please refer to [PKT 24.229], and [RFC 5626] for more information. If the protocol used for a SIP Session is UDP this value is used for SIP Timer D, otherwise is ignored and the SIP session.

- SubscribeRetry

  This attribute represents the retry period for the initial SUBSCRIBE due to error responses, the absence of a retry period in the Retry-After header response or a request timeout. Please refer to [PKT 24.229].

A.2.3.6 E164Mapping Object[24]

This object represents configured mapping of TNs to LAN ESEs. This is used when host-name in the Request-URI identifies an ESE WAN that is associated with multiple ESE LANs. In this case, the SBC consults the E164Mapping Object using the TN contained in the user-part of the Request-URI to identify the target ESE LAN.

Object Operations:

None

[24] ESG-N-12.0691-8 10/28/16 by PO

Table 8 - E164Mapping Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	Unsigned32	Key			
NumberPattern	SnmpAdminString	CRUD	SIZE(0..1024)		“”
LanESEName	AdminString	CRUD	SIZE(0..256)		

Attribute Descriptions:

- Id

  This key represents the unique identifier of an object instance.

- NumberPattern

  This attribute identifies the E.164 telephone numbers and abbreviated extension numbers that are mapped to the LAN ESE representing a SIP-PBX. The attribute consists of a comma-separated list of number entries, where each number entry is either an explicit number (i.e., an explicit E.164 number or an explicit abbreviated extension number) or an expression that identifies multiple numbers (e.g., a range of E.164 numbers or a range of abbreviated extension numbers).

  - A number entry that identifies an explicit E.164 number follows the global-number syntax of a tel URL as per [RFC 3966].
  - A number entry that identifies an explicit abbreviated extension number follows the local-number syntax of a tel URL defined in [RFC 3966] (i.e., a string of digits with no leading ‘+’ character).
  - A number entry that identifies multiple E.164 or extension numbers follows the E.164 or abbreviated extension syntax specified above, except that one or more digits are replaced with the wildcard character ‘X’ (e.g., the number entry ‘45XX’ identifies the range of extension numbers ‘4500’ to ‘4599’). When a received digit string matches multiple number entries, the best match is selected (e.g., for the case where there are two number entries ‘45XX’ and ‘456X’, the received digit string ‘4566’ matches the number entry ‘456X’).

  The ESG uses this attribute to route incoming calls received either in normal or survivable mode to the correct SIP-PBX.

  - LanESEName

  This attribute represents the LanESE Name associated with the NumberPattern.

A.2.3.7 QoS Object

This object contains the QoS parameters applicable to the ESG.

Object Operations:

None

Table 9 - QoS Object

Attribute Name	Type	Access	Type Constraints	Units	Default
RtpDSCP	unsignedByte	RU			0
RtcpDSCP	unsignedByte	RU			0
SipDSCP	unsignedByte	RU			0

Attribute Descriptions:

- RtpDSCP

  This attribute represents the DSCP value used for marking the 'valid' upstream RTP packets.

- RtcpDSCP

  This attribute represents the DSCP value used for marking the 'valid' upstream RTCP packets

- SipDSCP

  This attribute represents the DSCP value used for marking the 'valid' upstream SIP packets.

A.2.3.8 Servers Object

This object represents the list of servers that the ESG may need for proper operation.

Object Operations:

None

Table 10 - Servers Object

Attribute Name	Type	Access	Type Constraints	Units	Default
ENUMServer	InetAddress	RU			
LanDNS	InetAddress	RU			

Attribute Descriptions:

- ENUMServer

  This attribute represents the network address of the ENUM server in the enterprise network.

- LanDNS

  This attribute defines the network address of the DNS server in the enterprise network.

A.2.3.9 ProvALG object[25]

This object configures the mapping between the private LAN URL of the ESG and the public WAN URL of the HTTP Server. The LAN CPEs (aka ESEs) are configured with a Provisioning Server URL that matches the preconfigured LANUrl attribute of this object in the ESG. When the Provisioning ALG receives an HTTP request from a LAN CPE on the ESG LAN IP address:port associated with the ProvALG's preconfigured LANUrl attribute, it replaces the target LANUrl identified in the Request URI of the incoming request with the preconfigured WANUrl attribute of this object, and forwards the request to the Provisioning Server.

Object Operations

None

This is an abstract class.

Table 11 - ProvALG Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
LANUrl	Uri	CRUD			"
WANUrl	Uri	CRUD			"

Attribute Descriptions:

- Id

  This key represents the unique identifier of an object instance.

- LANUrl

  This attribute identifies the LAN URL used by a LAN CPE to send HTTP requests to the CPE's Provisioning Server. The LANUrl resolves to an IP address:port on the LAN interface of the ESG.

- WANUrl

  This attribute identifies the WAN URL of the Provisioning Server. The Provisioning ALG redirects incoming HTTP requests received from LAN CPEs and addressed to the LANUrl to this WANUrl.

A.2.3.9.1 CWMPALG object

This object is a specialization of the ProvALG for CWMP LAN/WAN mapping. When this object is instantiated, the Provisioning ALG behaves as a CWMP ALG as defined in section 8.3.1.

Object Operations

None

At least one instance is required to be created for this object

Attributes

[25] New sections added by ESG-N-12.0691-8 on 11/4/16 by PO

No additional attributes defined for this object.

A.2.3.9.2 HTTPALG object

This object is a specialization of the ProvALG for basic HTTP mapping. When this object is instantiated, the Provisioning ALG behaves as a HTTP ALG as defined in section 8.3.2.

Object Operations

None

At least one instance is required to be created for this object.

Attributes

No additional attributes defined for this object.

A.2.3.10 ProvLGStatus object

This object represents the current status of the Provisioning ALG WAN/LAN address mappings based on activity between the LAN CPE and the provisioning server.

Operations:

None

Table 12 - ProvALGStatus Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
LANUrl	Uri	R			
WANUrL	Uri	R			
CreationTime	DateTime	R			
LastUpdate	DateTime	R			

- Id

  This key represents the unique identifier of an object instance.

- CreationTime

  This attribute represents the date and time when the instance was created (e.g., when a newly rebooted CPE sent the first HTTP request to its provisioning server, or sent the first CWMP Inform to the ACS).

- LastUpdate

  This attribute represents the last time this instance was updated due to LAN CPE activity (e.g., when the CPE last sent an HTTP request to its provisioning server, or a CWMP Inform message to its ACS).

- LANUrl

  This attribute identifies the LAN URL used by a LAN CPE to send HTTP requests to the CPE's provisioning server. WANUrl

  This attribute identifies the WAN URL of the Provisioning Server.

- LANHosts

  This attribute represents a list of Hosts sharing the same ALG mapping. The list of host is normally a set of references to other objects describing the host information.

A.2.3.10.1 CWMPALGStatus object

This object represents the current status of the CWMP ALG WAN/LAN ConnectionRequestURL mappings.

The CWMP ALG maintains CWMP mappings based on the uniqueness of the LAN CPE device, which the ALGlearns from the CWMP inform Device Id. Therefore, only one host in the LANHosts attribute is reported per object instance.

Operations:

None

Table 13 - CWMPALGStatus Object

Attribute Name	Type	Access	Type Constraints	Units	Default
LANConnectionRequestUrl	Uri	R			
WANConnectionRequest	Uri	R			

- LANConnectionRequestUrl

  This attribute represents the ConnectionRequestURL the CWMP ALG learned from the CWMP Inform message received from the LAN CPE.

- WANConnectionRequestUrl

  This attribute represents the ConnectionRequestURL that the CWMP ALGassigns to a LANConnectionRequestUrl.

A.2.3.10.2 HTTPALGStatus object

This object represents the current status of the HTTP ALG mappings for HTTP requests received from the LAN CPE.

Object Operations

None

Attributes

No additional attributes defined for this object.

A.2.3.11 SBCMappingStatus Object

This object represents the status of valid and invalid WAN/LAN ESE bindings based on the WanESE and LanESE object associations. Invalid associations occur when the binding of WAN/LAN ESEs and/or the interworking function rule-set resolve to an invalid ESE or an exception, or when an error condition is detected at the time of the last processed message relate to the WAN/LAN ESE binding.

Object Operations:

None

Table 14 - SBCMappingStatus Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
SBCId	unsignedInt	key			
BindingTypes	Enum	R	signaling(1),media(2),mediaControl(3)		
WanESEName	AdminString	R	SIZE(0..256)		
WanESEAddress	InetAddress	R			
WanESEPort	InetPortNumber	R			
LanESEName	AdminString	R	SIZE(0..256)		
LanESEAddress	InetAddress	R			
LanESEPort	InetPortNumber	R			
SIPConnectMode	Enum	R	unknown(1),staticMode(2),registration Mode(3),		
Status	Enum	R	active(1),inactive(2),invalid(3)		
StatusDescription	AdminString	R	SIZE(0..256)		
LastTime	dateTime	R			
NextHopProxy	AdminString	R			
Release	boolean	RU			false

Attribute Descriptions:

- Id

  This key represents the unique identifier of an object instance.

- SBCId

  This attribute represents the Id of the SBC that provides this mapping.

- BindingTypes

  This attribute indicates the type of binding for the WAN/LAN ESE association.

- WanESEName

This attribute represents the WAN ESE Name of the mapping instance.

- WanESEAddress

This attribute represents the IP address of the UA, or empty if unknown.

- WanESEPort

This attribute represents the IP port that the ESE is using for this binding.

- LanESEName

This attribute represents the LAN ESE Name of the mapping instance.

- LanESEAddress

This attribute represents the IP address of the ESE.

- LanESEPort

This attribute represents the IP port that the ESE is using for this binding.

- SIPConnectMode

This attribute represents the SIPconnect 1.1 mode of operation of this mapping instance.

'unknown' indicates that the mode is undefined or unknown.

'staticMode' refers to SIPConnect 'static-mode'.

'registrationMode' refers to the SIP registration procedures.

- Status

This attribute represents the status of a WAN/LAN ESE mapping.

'active' indicates the mapping is valid and was successful.

'inactive' indicates the mapping is known to be valid but either WAN/LAN ESE or both are not operational.

'invalid' indicates the WAN/LAN ESE mapping was matched to an IWTemplate rule, but did not resolve in a known UE and/or ESE entity.

- StatusDescription

This attribute represents a human readable reason, or explanation of the current status of the SBC connection mapping.

- LastTime

This attribute represents the last time the status of a WAN/LAN ESE mapping was updated.

- NextHopProxy

This attribute represents the next hop (e.g., IBCF/P-CSCF).

- Release

This attribute when set to 'true' removes the binding from the SIP-NAT process. This action only applies to mappings of BindingType 'media' and 'mediaControl'. To update or remove 'signaling' type of bindings, the SBC ESEs need to be reconfigured (i.e., WanESE, LanESE). Reading this value always returns 'false'.

A.2.3.12 SBCFirewallLog Object

This object is used to store the log of events such as SIP Register violations and exceptions per the SBC SIP-aware firewall rule sets, including malformed and invalid SIP message occurrences detected by the SIP-aware firewall. This object may also contain violations related to promiscuous mode firewall operations, although these cases are normally reported by enterprise firewall models outside of the scope of this specification.

- Object Operations:

None

Table 15 - SBCFirewallLog Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	key			
SBCId	unsignedInt	key			
Direction	Enum	RU	inbound(1),outbound(2)		
DestinationEndPoint	AdminString	RU			
DestinationAddress	InetAddress	RU			
DestinationPort	InetPortNumber	RU			
SourceEndPoint	Uri	RU			
SourceAddress	InetAddress	RU			
SourcePort	InetPortNumber	RU			
StatusDescription	AdminString	RU			
Timestamp	dateTime	RU			

Attribute Descriptions:

- Id

This key represents the unique identifier of an object instance.

- SBCId

This attribute represents the Id of the SBC corresponding to the particular object instance. This attribute is used to further link the SBC and the LANESESIPFirewallUrl. For promiscuous mode firewall operations, this attribute is set to zero.

- Direction

This attribute represents the direction of the firewall exception.

'inbound' indicates messages received from the WAN facing interface,

'outbound' indicates messages received from the LAN facing interface.

- DestinationEndPoint

This attribute identifies the destination endpoint SIP URI of the log entry. For incoming packets (e.g., RTP, SIP) this attribute identifies the ESE. For outgoing packets this attribute identifies the remote endpoint.

- DestinationAddress

This attribute identifies the destination IP Address in the incoming and outgoing packets.

- DestinationPort

This attribute identifies the destination UDP/TCP port in the incoming and outgoing packets.

- SourceEndPoint

This attribute identifies the source endpoint SIP URI of the log entry. For incoming packets (e.g. RTP, SIP) this attribute identifies the remote endpoint. For outgoing packets this attribute identifies the ESE.

- SourceAddress

This attribute identifies the source IP Address in the incoming and outgoing packets.

- SourcePort

This attribute identifies the source UDP/TCP port in the incoming and outgoing packets.

- StatusDescription

This attribute provides human readable details on the reason this exception/log was generated.

- Timestamp

This attribute represents the time when this instance was logged.

A.2.3.13 SETA Object

This object represents the SIP endpoint Test Agent.

Object Operations:

None.

Table 16 - SETA Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Enable	boolean	CRUD			false
ESEReference	Uri	CRUD			
Ncr	unsignedInt	CRUD			10
PreferredCodec	AdminString	CRUD			PCMU
PacketizationTime	unsignedInt	CRUD		milliseconds	20
CallDuration	unsignedInt	CRUD		seconds	30

Attribute Descriptions:

- Enable

  When this attribute is set to 'true', SETA is enabled. When this attribute is set to 'false', SETA is disabled and any in-progress test calls are gracefully terminated.

- ESEReference

  This attribute identifies the User Agent to be used by SETA. Normally the ESEReference attribute will refer to an IMPU operating in the SETA 'role' (IMPU AdditionalInfo attribute with value 'ESG#SETA'). This SETA IMPU can be associated with the IMS Subscription of another IMPU. Additionally, the ESEReference attribute can refer to the name of an existing ESG Wan/LanESE, or the IMPU of an EDVA endpoint.

  If the ESEReference does not match an existing IMPU, and does not identify the ESE name of an existing Wan/LanESE, then it is treated as the name of a new instance of a Wan/LanESE object that is dedicated to the SETA function. In this case the Wan/LanESE is not associated with a SIP-PBX or SIP endpoint in the Enterprise network. This same name must be referenced by a SBCCfg instance in order to enable SETA to initiate and receive test calls through the SBC function.

- Ncr

  This attribute indicates the number of most recent calls to be captured in the SETA Call Statistics Log.

- PreferredCodec

  This attribute contains the Preferred network Codec List. The value in this object is formed as a comma-separated list of the well-known literal codec names in order of preference from left to right. The SETA must use the literal codec name as per RTP AV Profile [RFC 3551], or per encoding names registered with the IANA, or per encoding names referenced or defined in the PacketCable Codec-Media specification. Unknown or non-supported codecs are ignored. the zero-length string indicates the preferred codec list is vendor specific starting with G711 codecs.

- PacketizationTime

  This attribute represents the time duration in milliseconds of voice packetization.

- CallDuration

  This attribute represents the Call duration for the SETA end point to hang up. Calls may be graciously terminate by the other end.

A.2.3.14 SETAIncomingCfg Object

This object represents the configuration of SETA incoming call parameters.

Object Operations:

None

This object inherits all of its attributes from the base SETA object.

A.2.3.15 SETAOutgoingCfg Object

This object represents the configuration of outgoing call parameters for SETA.

Object Operations:

None

Table 17 - SETAOutgoingCfg Object

Attribute Name	Type	Access	Type Constraints	Units	Default
RemoteEndPoint	Uri	RU			
NATO	unsignedInt	RU			
Schedule	unsignedInt	RU		calls	0

Attribute Descriptions:

- RemoteEndPoint

  This attribute represents the called destination.

- NATO

  This attribute represents the no answer timeout used.

- Schedule

  This attribute represents the number of outgoing calls to be distributed over a 24 hours period. A value 0 indicates that when the SETA Enable attribute is set to 'true', only one outgoing call is initiated.

  Any time the value of the SETA Enable attribute is set to 'true' (even if previously set to 'true'), the SETA outgoing call schedule is restarted.

The remaining attributes are inherited from the base SETA object.

A.2.3.16 CallLogCtrl Object[26]

This object represents the control of the call statistics log. The handling of logs overruns when the log reaches the maximum size is vendor specific, and outside the scope of this specification.

Object Operations:

None.

Table 18 - CallLogCtrl Object

Attribute Name	Type	Access	Type Constraints	Units	Default
CallType	Enum	RU	sbc(1), seta(2),all(3),none(4)		none
Enable	boolean	RU			False
Filter	AdminString	RU	SIZE(0..1024)		

[26] ESG-N-12.0691-8 10/28/16 by PO

- CallType

  This attribute represents the type of call to be logged

  'sbc' indicates calls traversing the sbcs are logged.

  'seta' indicates SETA calls are logged.

  'all' indicates both SBC and SETA calls are logged

  'none' indicates no calls will be logged.

- Enable

  This attribute controls whether calls are logged or not. When set to 'true', call logging is started. When set to 'false', call logging is stopped.

- Filter

  This attribute represents an expression to identify a subset of calls to be logged. The Filter expressions may follow the syntax of common packet sniffing tools, or may be vendor-specific.

A.2.3.17 CallLog Object

This object represents common attributes for call logs. As an abstract class this object is not instantiated directly, but inherited by other objects.

Object Operations:

None.

Table 19 - CallLog Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Id	unsignedInt	Key			
StartTime	dateTime	R			
EndTime	dateTime	R			
CallID	unsignedInt	R			
Direction	Enum	R	inbound(1),outbound(2)		
OriginatingSIPURI	Uri	R			
TerminatingSIPURI	Uri	R			

Attribute Descriptions:

- Id

  This key represents the unique identifier of the instance. This key is a monotonically increasing integer number. When the buffer is full the lowest index instance is deleted to allow writing the newest record.

- StartTime

  This attribute represents the time the call starts.

- EndTime

  This attribute represents the time the call ends.

- CallID

  This attribute represents the SIP Call-ID of the call.

- Direction

  This attribute represents the direction of the call with respect to the measuring device.

- OriginatingSIPURI

  This attribute represents the originating SIP URI of the call.

- TerminatingSIPUR

  This attribute represents the terminating SIP URI of the call.

A.2.3.18 SBCCallStats Object

This object represents the call statistics log of SBC calls

Object Operations:

None

All attributes are inherited from the CallLog object

A.2.3.19 SETACallStats Object

This object contains the statistics of calls traversing the ESG.

Object Operations:

None.

Table 20 - SETACallStats Object

Attribute Name	Type	Access	Type Constraints	Units	Default
SRD	unsignedInt	R		seconds	
SDD	unsignedInt	R		seconds	
ResponseCode	unsignedInt	R			
CallCompletionStatus	Enum	R	success(1),fail(2)		

Attribute Descriptions:

- SRD
  This attribute corresponds to Session Request Delay (SRD).

- SDD

  This attribute corresponds to Session Disconnect Delay (SDD).

- ResponseCode

  This attribute corresponds to the SIP response code of the call.

- CallCompletionStatus

  This attribute corresponds to the call completion status. Possible values are:

  'success': The call was answered and terminated as expected.

  'fail': The call was not completed successfully.

The remaining attributes are inherited from the CallLog object

A.2.3.20 TraceLogCtrl Object[27]

This object controls the collection and upload of call SIP and RTP traces.

Object Operations:

None

Table 21 - TraceLogCtrl Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Enable	boolean	RU			
CallType	Enum	RU	sbc(1), seta(2),all(3),none(4)		none
Upload	Uri	RU			
LastUploadTime	dateTime	RU			
LastUploadStatus	AdminString	RU	success(1),failed(2), inprogress(3),none(4)		none
Filter	AdminString	RU	SIZE(0..1024)		

Attribute Descriptions:

- Enable

  This attribute controls the start and end of call traces collection. The value 'true' enables the collection of traces into the traces file. The value 'false' stops collection of traces. If the value 'true' is set while a log file is being currently uploaded, the ESG will always enable the collection of traces immediately, and may either abort or continue the upload process.

- CallType

  This attribute represents the type of call to be logged.

  'sbc' indicates calls traversing the sbcs are logged.

  'seta' indicates SETA calls are logged.

[27] ESG-N-12.0691-8 10/28/16 by PO

'all' indicates both SBC and SETA calls are logged.

'none' indicates no calls will be logged.

- Upload

  This attribute represents the FTP URL where the trace log is uploaded. The FTP URL follows [RFC 3986] recommendation. When set to a value the trace file is immediately uploaded.

- LastUploadTime

  This attribute indicates the last time the trace log was attempted to be uploaded.

- LastUploadStatus

  This attribute indicated the status of the last attempt to upload the trace log.

  'success' indicates the file upload was completed.

  'failed' indicates the file upload and retry failed.

  'inProgress' indicates the file is in the upload process.

  'none' indicates no upload activity is being performed

- Filter

  This attribute represents an expression to identify a subset of packets to be logged in the trace capture. The Filter expressions may follow the syntax of common packet sniffing tools, or may be vendor-specific.

A.2.3.21 SIPTraceLogCrl Object

This object controls the SIP trace logging.

Object Operations:

None.

Table 22 - SIPTraceLogCrl Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Ncr	unsignedInt	RU		calls	10

Attribute Descriptions:

- Ncr

  This attribute indicates the number of calls to log after the trace log is enabled.

A.2.3.22 RTPTraceLogCtrl Object

This object controls the RTP trace logging.

Object Operations:

None.

Table 23 - RTPTraceLogCtr Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Tcr	unsignedInt	RU		seconds	0

Attribute Descriptions:

- Tcr

  This attribute indicates the number of seconds to log after the trace log is enabled.

A.2.3.23 PublishCtrl Object[28]

This object represents the attributes used to control the SIP PUBLISH messages.

Object Operations:

None

Table 24 - PublishCtrl Object

Attribute Name	Type	Access	Type Constraints	Units	Default
Enable	boolean	RU			false
DebugMode	boolean	RU			false
UA	AdminString	RU	SIZE(0..256)		
Collector1	InetAddress	RU			
Collector2	InetAddress	RU			
WanESEList	AdminString	RU			

Attribute Descriptions:

- Enable

  This attribute is used to enable or disable the sending of SIP PUBLISH messages to the Telemetry collector. If the value of this attribute is set to true, the ESG is enabled to send SIP PUBLISH messages.

- DebugMode

  The section 6.4.1.1.4 explains three types of metrics report: session reports, interval reports and alert reports. If the SIP PUBLISH mechanism is enabled (by setting the value of enable attribute to true) and the value of this attribute is set to false, then the ESG MUST only send the session reports. If the value of both the enable and DebugMode is set to true, then the ESG MUST send all three types of reports, if all three types of reports are supported. If the ESG does not support all three types of report, then the ESG MUST throw a configuration error when the user tries to set the value of DebugMode to true.

[28] ESG-N-12.0691-8 10/28/16 by PO

- UA

  This attribute represents the UA that sends the PUBLISH messages. This UA could be a textual reference to a configured IMPU in the ESG, or a separate UA.

  As an exception to PacketCable Users configuration, the PUBLISH UA IMPU being referenced may be associated with an IMPI and/or a SIP network (domain).

  If no IMPI is associated, no authentication is performed during the UA registration.

  If no network is associated with the IMPU, the UA does not register with a SIP Core and perform publishing directly to the collector. In this case if the IMPU referenced contains an IMPI, authentication may be performed on a message-by-message case, based on collector behavior.

- Collector1

  This attribute contains the address of the primary Telemetry collector. The UA sends the SIP PUBLISH message to this address.

- Collector2

  This attribute contains the address of the secondary Telemetry collector. The UA sends the SIP PUBLISH message to this address, if the primary telemetry server is down or out of service.

- WanESEList

  This attribute contains a comma-separated list of WanESE names. The SIP PUBLISH messages for these WanESEs will be sent to the collectors using the UA configured in this instance.

  If the UA and collector (1 and 2) attributes are not configured, then the SIP PUBLISH messages are sent to the next hop proxy for each WanESE on the list.

Appendix I Acknowledgements[29]

We wish to thank the vendor participants contributing directly to this document:

PacketCable wishes to recognize the following individuals for their significant involvement and contributions to this specification (ordered alphabetically by company name and individual's first names in each company):

For development of the initial I01 version of the document:

Peter Som De Cerff, Adtran

Doug Wadkins, Edgewater

Eugene Nechamkin, Broadcom

Gordon Li, Broadcom

Guhan Parthasarathy, Global Edge Software Ltd.

Lee Valerius, Huawei

Harprit Chhatwal, InnoMedia

Geoff Devine, SMC Networks

Satish Kumar, Texas Instruments

Eduardo Cardona and Vikas Sarawat, CableLabs staff members, are thanked for their direct contributions to this specification.

For development of the final version of the document:

David Schenkel, Adtran

Bob Hornburg, Audiocodes,

Harprit Chhatwal, Innomedia

Eduardo Cardona, TWC

David Hancock and the Business Services Team (CableLabs)

[29] ESG-N-12.0691-8 on 11/4/16 by PO

Appendix II Revision History

The following Engineering Changes have been incorporated in PKT-SP-ESG-C01-170405.

ECN Identifier	Accepted Date	Title of EC	Author
ESG-N-12.0691-8	11/5/12	ESG Provisioning-aware NAT/Firewall	Hancock
ESG-N-12.0682-8	7/16/12	ESG Survivability and Data Model Updates	Hancock
ESG-N-11.0665-6	7/11/11	Mandate ESG Support of IPv4/6 Dual-Stack	Hancock