

PacketCable™

Enterprise SIP Gateway Technical Report

PKT-TR-ESG-V01-101103

RELEASED

Notice

This PacketCable technical report 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. 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.

Neither CableLabs nor any member company is responsible to any party for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this document, or any document referenced herein. 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.

© Copyright 2010 Cable Television Laboratories, Inc.

All rights reserved.

Document Status Sheet

Document Control Number:	PKT-TR-ESG-V01-101103			
Document Title:	Enterprise SIP Gateway Technical Report			
Revision History:	V01 –11/02/10			
Date:	November 3, 2010			
Status:	Work in Progress	Draft	Released	Closed
Distribution Restrictions:	Author Only	CL/Member	CL/Member/Vendor	Public

Trademarks

Advanced Digital Cable™, CableCARD™, CableHome®, CableLabs®, CableNET®, CableOffice™, CablePC™, DCAS™, DOCSIS®, DPoE™, EBIF™, eDOCSIS™, EuroDOCSIS™, EuroPacketCable™, Go2Broadband™, M-Card™, M-CMTS™, OCAP™, OpenCable™, PacketCable™, PCMM™, and tru2way® are marks of Cable Television Laboratories, Inc. All other marks are the property of their respective owners.

Contents

1	SCOPE	1
1.1	Introduction and Purpose	1
1.2	Overview	1
1.2.1	<i>Embedded ESG</i>	2
1.2.2	<i>Stand-Alone ESG</i>	3
2	REFERENCES	4
2.1	Normative References	4
2.2	Informative References	4
2.3	Reference Acquisition	5
3	TERMS AND DEFINITIONS	5
4	ABBREVIATIONS AND ACRONYMS	7
5	ESG FUNCTIONAL REQUIREMENTS	9
5.1	General Device Requirements	9
5.2	SIP Header Manipulation Requirements	9
5.3	SIP-ALG Security Requirements	10
5.4	Network Address Translation	10
5.5	Firewall Requirements	10
5.6	Media Relay Requirements	11
5.7	Quality of Service	11
5.8	Operations, Administration, Management, and Provisioning	11
5.9	Event Logging, Voice Statistics, and Reporting	12
5.10	Embedded SIP Endpoint Test Agent	12
5.11	Availability and Reliability	13
6	ESG ARCHITECTURE	14
6.1	Session Border Controller	15
6.1.1	<i>SIP Interworking Function</i>	15
6.1.2	<i>SIP-aware NAT/Firewall Functions</i>	18
6.2	SETA Function	19
6.2.1	<i>Base Functionality</i>	19
6.2.2	<i>RTP Loopback</i>	19
6.2.3	<i>Test Call Origination and Termination</i>	20
6.3	Telemetry Function	20
6.3.1	<i>Voice Statistics</i>	20
6.3.2	<i>SIP and RTP Trace</i>	22
6.4	Quality of Service	22
6.5	Security Services	23
6.5.1	<i>SIP Signaling Security</i>	23
6.5.2	<i>Security Services for RTP Media</i>	23
6.5.3	<i>Security Services for Provisioning & Management Interfaces</i>	23
6.6	Operations, Administration, Management, and Provisioning	24
APPENDIX I	ACKNOWLEDGEMENTS	25

Figures

Figure 1 - ESG Functional Diagram.....2
Figure 2 - Internal Architecture of ESG14
Figure 3 - Block Diagram of SBC acting as B2BUA.....17
Figure 4 - SBC as NAT only17
Figure 5 - Block Diagram of SBC acting as transparent pipe18

1 SCOPE

This technical report describes the functional requirements and technical solution for a new PacketCable™ 2.0 component called the Enterprise SIP Gateway (ESG). The ESG is deployed as a gateway device at the demarcation point between a Service Provider and Enterprise network to facilitate the delivery of Business Voice services to enterprise customers, including SIP Trunking service to an enterprise SIP-PBX, and hosted IP Centrex service to enterprise SIP phones.

The information provided in this document is informative in nature, and serves as input to the normative ESG requirements defined in [PKT-SP-ESG].

1.1 Introduction and Purpose

The ESG provides a number of advantages to MSOs deploying Business Voice services to enterprise customers:

- It improves the scalability of deployment. Due to a lack of industry standards in this area, there is a large divergence among the SIP implementations supported by the various makes and models of Business Voice CPE equipment deployed in today's enterprise networks. As a result, operators that want to deploy Business Voice must spend a lot of time and energy testing for and resolving interoperability issues between their own PacketCable network and the enterprise SIP-PBX and SIP endpoint devices. Also, once the service is deployed, any version/software update of the enterprise SIP devices can create new interworking issues. This can become an excessive burden to operators as the number of Business Voice deployments increases. The ESG mitigates this problem by normalizing the SIP signaling across the wide variety of Business Voice CPE equipment to a common standard that is compatible with the Service Provider network.
- It improves security. The introduction of Business Voice creates a new deployment model where stand-alone and non-certified endpoint devices are attached to the PacketCable network. In the case of PacketCable residential telephony services, the E-DVA devices are typically certified to comply with PacketCable specifications, and therefore, can be counted on to behave in a manner compatible with the Service Provider network. Also, because the SIP endpoint in the E-DVA is embedded, the operator can be confident that it won't be tampered with and start behaving in a manner that is harmful to the network. Stand-alone CPE devices deployed in enterprise networks don't have either of these protections. The ESG resolves this issue by implementing a SIP-aware firewall that can protect the cable operator's network from improper or malicious CPE behavior.
- It simplifies fault identification, isolation, and recovery. The ESG supports proactive test tools and data monitoring capabilities that enable MSOs to more quickly and efficiently detect, locate, and resolve component failures and network problems that can block the delivery of Business Voice services to the enterprise. This reduces the time and cost to the MSO in performing this function, and increases the level of service to the enterprise customer.

1.2 Overview

The Enterprise SIP Gateway device comes in two versions; an embedded version, where the ESG contains a DOCSIS version 1.1 or later cable modem, and a stand-alone version where the ESG is connected to the access network via a standard Ethernet interface. Figure 1 shows a high-level block diagram of the embedded version. The ESG sits at the boundary between the Service Provider and Enterprise network, and acts as a well-defined demarcation point for the business voice traffic that flows between the two networks. The SIP signaling interface on the right-hand-side of the ESG toward the Enterprise network does not comply with a specific standard, but varies to interwork with the specific CPE device deployed in the enterprise network. The SIP signaling interface on the left-hand-side between the ESG and the Service Provider network conforms to PacketCable specifications. The Provisioning and Management interfaces connect the ESG to the operator's various back-office OSS systems to enable the Service Provider to configure, control and monitor the ESG.

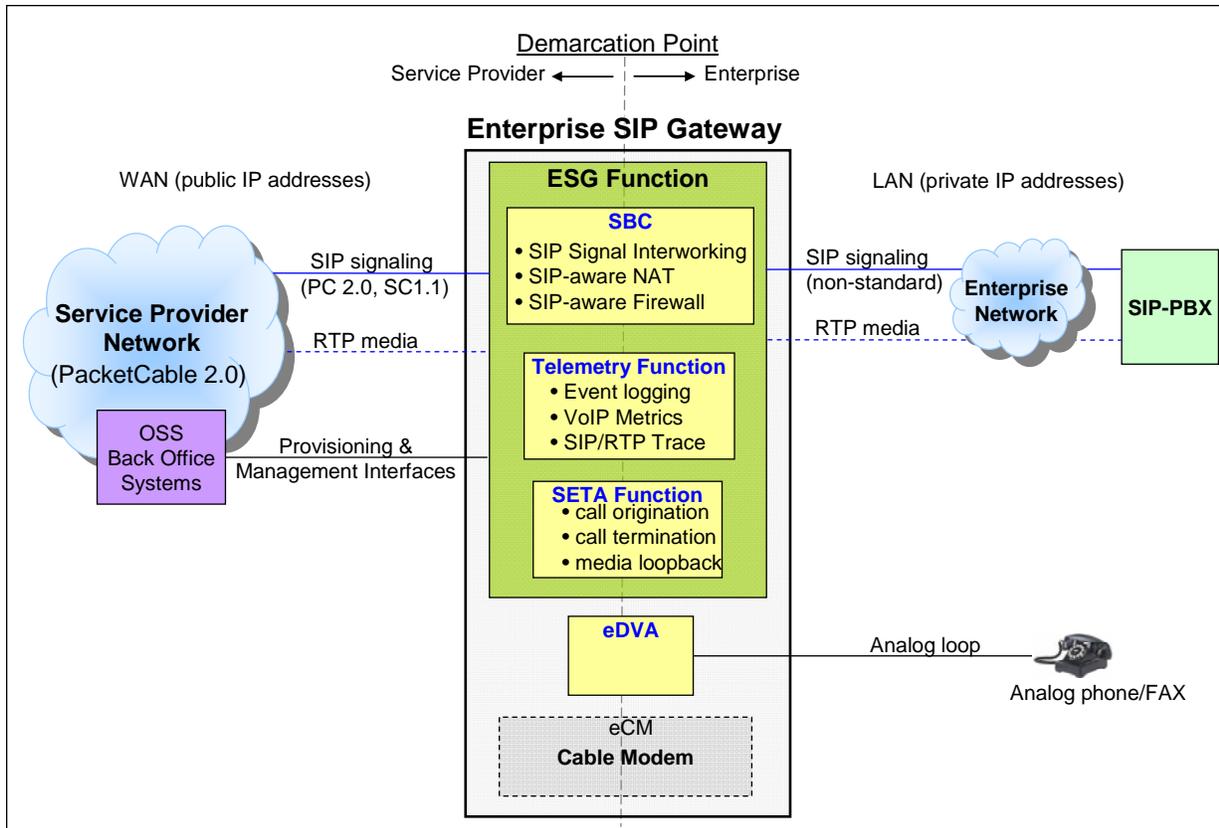


Figure 1 - ESG Functional Diagram

1.2.1 Embedded ESG

As shown in Figure 1, the embedded version of the ESG device contains three major blocks; the ESG Function which is the focus of this Technical Report, an embedded DVA serving analog lines, and the embedded cable modem.

The ESG Function itself contains three sub-functions:

- A Session Border Controller (SBC) that performs signal interworking to normalize the non-standard SIP signaling procedures supported by the enterprise CPE devices into a standard version of SIP that is compatible with the PacketCable 2.0 network. The SBC also contains a SIP-aware NAT that translates the IP addresses contained in the IP headers and SIP messages between the enterprise private LAN addresses and the Service Provider public WAN addresses. Finally, the SBC contains a SIP-aware firewall that enforces operator-configured policies for SIP messages entering the Service Provider network, to ensure that the messages are well-formed, from valid sources, etc.
- A Telemetry Function that collects and reports data associated with business voice traffic, such as VoIP Metrics, error event logs, and SIP and RTP traces.
- A SIP Endpoint Test Agent (SETA) that can 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).

The eDVA provides analog line (FXS) service to support enterprise FAX service, and to provide a reliable telephony connection for reporting alarms from an enterprise alarm panels.

The embedded eCM provides high-speed data connectivity to the PacketCable core network over DOCSIS/HFC access network.

1.2.2 Stand-Alone ESG

The stand-alone version of the ESG contains only the ESG Function shown in Figure 1 (i.e., it does not contain an embedded eCM or eDVA). The stand-alone version supports a standard RJ45 Ethernet port as its WAN interface to the PacketCable network.

2 REFERENCES

2.1 Normative References

There are no normative references in this document.

2.2 Informative References

This technical report uses the following informative references.

- [ID-SIP-RTCP] IETF Internet Draft, Session Initiation Protocol Event Package for Voice Quality Reporting, draft-ietf-sipping-rtcp-summary-13, August 4, 2010.
- [MULPI3.0] DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-I14-101008, October 8, 2010, Cable Television Laboratories, Inc.
- [OSSIV2] Operations Support System Interface Specification, CM-SP-OSSIV2.0-C01-081104, November 4, 2008, Cable Television Laboratories, Inc.
- [OSSIV3] DOCSIS 3.0 Operations Support System Interface Specification, CM-SP-OSSIV3.0-I13-101008, October 8, 2010, Cable Television Laboratories, Inc.
- [PCMM] PacketCable Multimedia Specification, PKT-SP-MM-I05-091029, October 29, 2009, Cable Television Laboratories, Inc.
- [PKT-EUE-DATA] PacketCable E-UE Provisioning Data Model Specification, PKT-SP-EUE-DATA-I05-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [PKT-EUE-PROV] PacketCable E-UE Provisioning Framework Specification, PKT-SP-EUE-PROV-I05-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [PKT-RST-E-DVA] PacketCable Residential SIP Telephony E-DVA Specification, PKT-SP-RST-E-DVA-I07-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [PKT-RST-EUE-PROV] RST E-UE Provisioning Specification, PKT-SP-RST-EUE-PROV-I05-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [PKT-SP-BSS] PacketCable Business SIP Services (BSS) Feature Specification, PKT-SP-BSSF-I03-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [PKT-SP-ESG] Enterprise SIP Gateway Specification, PKT-SP-ESG-I01-101103, November 3, 2010, Cable Television Laboratories, Inc.
- [PKT-TR-SEC] PacketCable Security Technical Report, PKT-TR-SEC-V05-080425, April 25, 2008, Cable Television Laboratories, Inc.
- [PKT-UE-DATA] PacketCable 2.0 UE Provisioning Data Model, PKT-SP-UE-DATA-I02-100527, May 27, 2010, Cable Television Laboratories, Inc.
- [RFC 3261] IETF RFC 3261, SIP: Session Initiation Protocol, June 2002.
- [RFC 3264] IETF 3264, An Offer/Answer Model with Session Description Protocol (SDP), June 2002.

- [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 5424] IETF RFC 5424, The Syslog Protocol, March 2009.
- [SIP-CONNECT1.1] SIP-PBX/Service Provider Interoperability, SIPconnect 1.1 Technical Recommendation, Draft, SIP Forum, 2009.

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>
- Internet Engineering Task Force (IETF) Secretariat, 46000 Center Oak Plaza, Sterling, VA 20166, Phone +1-571-434-3500, Fax +1-571-434-3535, <http://www.ietf.org>
- SIP Forum, Internet: <http://sipforum.com/>

3 TERMS AND DEFINITIONS

This specification uses the following terms:

Back-to-Back User Agent	A back-to-back user agent (B2BUA) is a logical entity defined in [RFC 3261], [RFC 3261], [RFC 3261], [RFC 3261], and [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 its 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.
Enterprise SIP Entity	A SIP-PBX or a SIP endpoint, located in the Enterprise 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.
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 Trunking Service	The Business Voice deployment model where the Service Provider network provides network connectivity to a SIP-PBX located in the Enterprise network.
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.

4 ABBREVIATIONS AND ACRONYMS

This document uses the following abbreviations:

ALG	Application Layer Gateway
ARP	Address Resolution Protocol
ASN.1	Abstract Syntax Notation One
B2BUA	Back-to-Back User Agent
CM	Cable Modem
CPE	Customer Premises Equipment
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
eSAFE	Embedded Service/Application Functional Entity
ESE	Enterprise SIP Entity
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
MAC	Media Access Control
MIB	Management Information Base
NAPT	Network Address Port Translation
NAT	Network Address Translation
OID	Object ID
QoS	Quality of Service
PC 2.0	PacketCable 2.0
RFC	Request For Comment
SDP	Session Description Protocol
SETA	SIP Endpoint Test Agent
SIP	Session Initiation Protocol
SP	Service Provider
SP-SSE	Service Provider SIP Signaling Entity
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security

TLV	Type/Length/Value
TTL	Time To Live
UDP	User Datagram Protocol
UE	User Equipment
VoIP	Voice over Internet Protocol
WAN	Wide Area Network

5 ESG FUNCTIONAL REQUIREMENTS

This section identifies the high-level functional requirements that must be supported by the Enterprise SIP Gateway. These functional requirements are supported by the normative ESG requirements defined in [PKT-SP-ESG].

5.1 General Device Requirements

The PacketCable ESG (referred to as the "ESG" going forward) can be deployed in one of two versions; an embedded version where the ESG contains an embedded DOCSIS version 1.1 or later cable modem with an RF coaxial interface, and a stand-alone version where the ESG is separate from the cable network and connected to it via a standard RJ45 10/100/1000BaseT Ethernet port.

The embedded version of the ESG must contain an embedded PacketCable 2.0 E-DVA with four (4) RJ-11 ports. The stand-alone ESG may support analog lines as an implementation option, either embedded in the ESG itself or as a stand-alone SIP-to-analog-line terminal adapter. The analog line for the stand-alone ESG is not required to support the already-defined PacketCable 2.0 S-DVA provisioning procedures in the initial release (this may be required in a future release).

The ESG must provide at least two (2) 10/100/1000BaseT Ethernet LAN ports.

1. One port dedicated to the delivery of Business Voice service.
2. One port dedicated to the delivery of High Speed Data (HSD) service.

The ESG may also support a configuration option where both voice and HSD service are provided on a single LAN port.

The ESG must be capable of supporting a minimum of 24 simultaneous call sessions.

The ESG must support IPv6 for all network interfaces, applications, and tools.

The ESG must support both SIP-PBX and hosted IP-Centrex service.

5.2 SIP Header Manipulation Requirements

The ESG must provide configuration controls that enable the MSO to define header manipulation rules that normalize the SIP signaling procedures from the enterprise SIP entity to standard SIP procedures defined by PacketCable toward the Service Provider network. The PacketCable standard is SIPconnect1.1 for SIP-PBX SIP Trunking service, and PacketCable 2.0 BSS for hosted IP Centrex service.

The ESG must provide both course-grained and fine-grained configuration controls (must not be mutually exclusive):

- course-grained where the MSO selects the interworking instance (e.g., selects the SIP-PBX make/model), and
- fine-grained where the MSO programs header manipulation rules.

The ESG must be capable of replacing an IP address with FQDN (Fully qualified domain name) and vice versa.

5.3 SIP-ALG Security Requirements

The ESG must support the ability to establish a TLS connection with the MSO network to secure business voice SIP messages, with the ESG playing the role of TLS client. The ESG must support the ability to enable or disable the TLS connection, based on local configuration. For the embedded ESG, the TLS connection for business voice traffic is in addition to, and can be enabled/disabled independently of, the E-DVA TLS connection. When TLS is enabled, the ESG must ensure that all SIP messages associated with business voice traverse the TLS connection. The ESG must support the client requirements for both TLS server and TLS client authentication.

The ESG must support a configuration option to force all SIP signaling messages to be sent via TCP or UDP. The ESG can interwork between TCP & UDP between its WAN and LAN interfaces.

The ESG must support SIP digest authentication as defined in PacketCable 24.229 and SIPConnect 1.1. The ESG must play the role of digest client, responding to 401-Unauthenticated and 407-Proxy-Authentication-Required challenges from the MSO network. The ESG must support a configuration option that enables the MSO to disable Digest authentication.

5.4 Network Address Translation

The ESG must provide network IP address and port address traversal between the Enterprise and Service Provider networks for both SIP messaging and RTP traffic.

The ESG must ensure that on its WAN interface, the RTP for a session is an even port, and the RTCP port for the same session is the RTP port+1.

ESG must maintain NAT address bindings for SIP signaling address of all directly addressable entities in enterprise network, for example:

- the registered contact address of the SIP-PBX,
- the registered contact address of each enterprise SIP endpoint for hosted IP-Centex,
- the contact address of enterprise endpoints advertised in a dialog-initiating request/response.

The ESG must perform IPv4/6 interworking (e.g., to enable a SP network on IPv6 to serve an enterprise on IPv4).

5.5 Firewall Requirements

The ESG must provide a SIP-aware firewall with an option to enable and disable the firewall function.

The ESG should provide the capability to relay or block message bodies based upon inspection of Content-Type header.

The ESG should provide the capability to relay or block message bodies in multipart mime-body SIP requests.

The ESG should provide the ability to configure rate-limiting on a per-host, per-user, per-network, per-message type, and per-device basis.

The ESG should provide the ability to filter (drop/permit) traffic based on a discretionary number of parameters, including but not limited to: source IP address, destination IP address, source TCP/UDP port, destination TCP/UDP port, source user, destination user, and service. The ESG must allow the following relational combinations of parameters: AND, OR, EXCLUSIVE OR, NOT.

The ESG should provide the ability to filter for various SIP Methods such as ACK, BYE, CANCEL, INFO, INVITE, MESSAGE, NOTIFY, OPTIONS, PRACK, PUBLISH, REFER, REGISTER, SUBSCRIBE and UPDATE.

The ESG must provide a mechanism that enables the Service Provider to configure the default behavior of the filter to drop all traffic not explicitly allowed (white-list) or to selectively drop specified traffic (blacklist) based on arbitrary parameters.

The ESG must log information related to invalid/discarded messages (e.g., what was discarded, why it was discarded). At a minimum, the number of discarded messages must be logged.

The ESG must be able to dynamically open ports for media transit based on SDP information. The ESG must support dynamic port allocation based on the SDP Offer/Answer exchanges [RFC 3264].

5.6 Media Relay Requirements

The PacketCable ESG must be transparent to RTP media encoding.

The ESG should support multiple RTP sessions per SIP dialog. For example, a SIP dialog may establish an RTP session for audio and a second RTP session for video.

5.7 Quality of Service

The ESG must be capable of enforcing bandwidth control policies to prioritize SIP and RTP packets over other traffic.

The ESG must be capable of independently setting DiffServ Code Points (DSCP) for signaling and media packets as well as provisioning flows.

The ESG must be capable of separating management network traffic from all other network traffic.

The ESG must ensure that upstream "best effort" high-speed-data does not impair upstream voice traffic (SIP signaling and RTP) within the ESG before it gets on the DOCSIS QoS service flow.

5.8 Operations, Administration, Management, and Provisioning

For the embedded ESG, Cable modem provisioning must comply with DOCSIS operations support system specifications [OSSIV2] and [OSSIV3].

For the embedded ESG, the E-DVA provisioning must comply with PacketCable 2.0 [PKT-EUE-PROV].

The ESG may support SNMPv3 and SNMPv2 (SNMPv2 'compatibility mode') [OSSIV2], and [OSSIV3].

The ESG may support SSHv2 or higher.

The ESG should provide a web GUI with secure HTTPS access for device configuration.

The ESG should support NTP for time synchronization.

The ESG should store all configuration data associated with business voice service in non-volatile memory (i.e., the data must survive power cycles).

ESG should support maintenance commands that enable Service Provider to remove the ESG from service. The ESG should support both a "graceful" remove-from-service command, where active calls are allowed to drain before the ESG is removed from service, and "forced" remove-from-service command where the ESG releases all active calls and transitions to out-of-service state immediately. The ESG should support configuration controls that enable the Service Provider to set the error response code or local announcement that is provided to new-call requests from the enterprise when the ESG is out-of-service.

5.9 Event Logging, Voice Statistics, and Reporting

The ESG must measure, collect, and store VoIP performance metrics per call, including: Incoming or Outgoing, Originating and Terminating Phone Numbers, Call Timestamp, Call Disposition, Call Duration, Call History, R-Factor, packet loss, jitter and latency. The ESG will collect and store two sets of metrics per call; one for each incoming RTP stream.

The ESG must store collected metrics and performance statistics in a local log file.

The ESG must use Syslog as defined in [RFC 5424] to make information available to the service provider.

The ESG must provide direct SNMP access to IP packet performance statistics, VoIP call logs, and quality statistics.

The ESG should provide SNMP configurable alarm thresholds on QoS parameters that will generate an SNMP trap.

The ESG must capture and store SIP signaling traces per call for at least the last 10 calls. The traces will show both the ingress and egress form of each SIP message that traverses the ESG.

The ESG must be able to capture and store RTP media for calls traversing the media relay. This functionality will be used for diagnostics, audio playback, and troubleshooting purposes. Additionally, it should provide the ability to filter the traffic capture based on signaling packets, RTP packets for a particular SIP URI or IP address.

The ESG should provide the ability to run a packet capture of TCP/IP traffic traversing the LAN interface with options to filter the traffic capture based on all traffic, signaling only, signaling and RTP, or traffic for a specific SIP URI or IP address.

The Service Provider should be able to configure trace start/stop criteria (e.g., trace all SIP messages in circular buffer until SIP error response 421 received).

5.10 Embedded SIP Endpoint Test Agent

This section presents the requirement for a SIP Endpoint Test Agent (SETA).

The PacketCable ESG must implement a SIP Endpoint Test Agent (SETA).

The SETA must originate and terminate calls under program control.

The SETA must support manually initiated and scheduled test calls under program control.

The SETA must originate test calls according to an MSO-configured set of parameters for dialed numbers, calling frequency, preferred CODEC/packetization rate/etc., and call duration with RTP packet generation.

The SETA must have the ability to load and play a specific audio file. The audio file is to be replayed over and over for the call duration. This enables collection of VoIP metrics at remote endpoint.

The SETA must report the call disposition of test calls (success, call setup time, post dial delay, busy, blocked, dropped, etc.) and call quality (loss, latency, jitter, etc.).

The SETA must provide the ability to define alarm thresholds and generate SNMP traps on call disposition with SIP error code and call quality (e.g., raise alarm if periodic test call fails). The Service Provider should be able to control whether or not SETA failures are reported to network, and if reported, whether as an event or an alarm, and how they're reported (SNMP vs. Syslog).

The ESG must support RTP loopback at the media relay following the procedures defined in [PKT-RST-E-DVA]. Support for media loopback is not required. Also, the media-relay RTP loopback test should be independent of the E-DVA loopback tests, i.e., it should not require the use of any E-DVA resources, or affect the ability of the E-DVA to establish calls.

5.11 Availability and Reliability

The ESG can provide an internal battery that will power all functions with normal operation of the ESG for two hours during a customer premise power failure.

If internal battery is provided, and customer premise power is not restored after two hours, the ESG must disable the Ethernet LAN interface and provide six hours minimum of remaining battery capacity to power the cable modem and E-DVA RJ-11 ports.

The ESG should support an external battery.

The ESG must comply with GR-1089 standards for devices attached to inside wiring.

6 ESG ARCHITECTURE

Figure 2 shows the internal architecture of the ESG. As shown in the diagram, the SBC within ESG provides the only conduit between the enterprise Local-Area-Network (LAN) and the Service Provider Wide-Area-Network (WAN) in order to carry business voice traffic between the SIP entities in the enterprise 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 between the enterprise SIP entities (PBX and hosted endpoints) and the LAN side of the SBC. Interfaces (3) and (4) in turn carry the SIP signaling and RTP media respectively between the WAN side of the SBC and the PacketCable 2.0 Network.

Note: This diagram is for illustrative purposes only as a means of describing the behavior of the ESG, and is not meant to mandate a specific implementation.

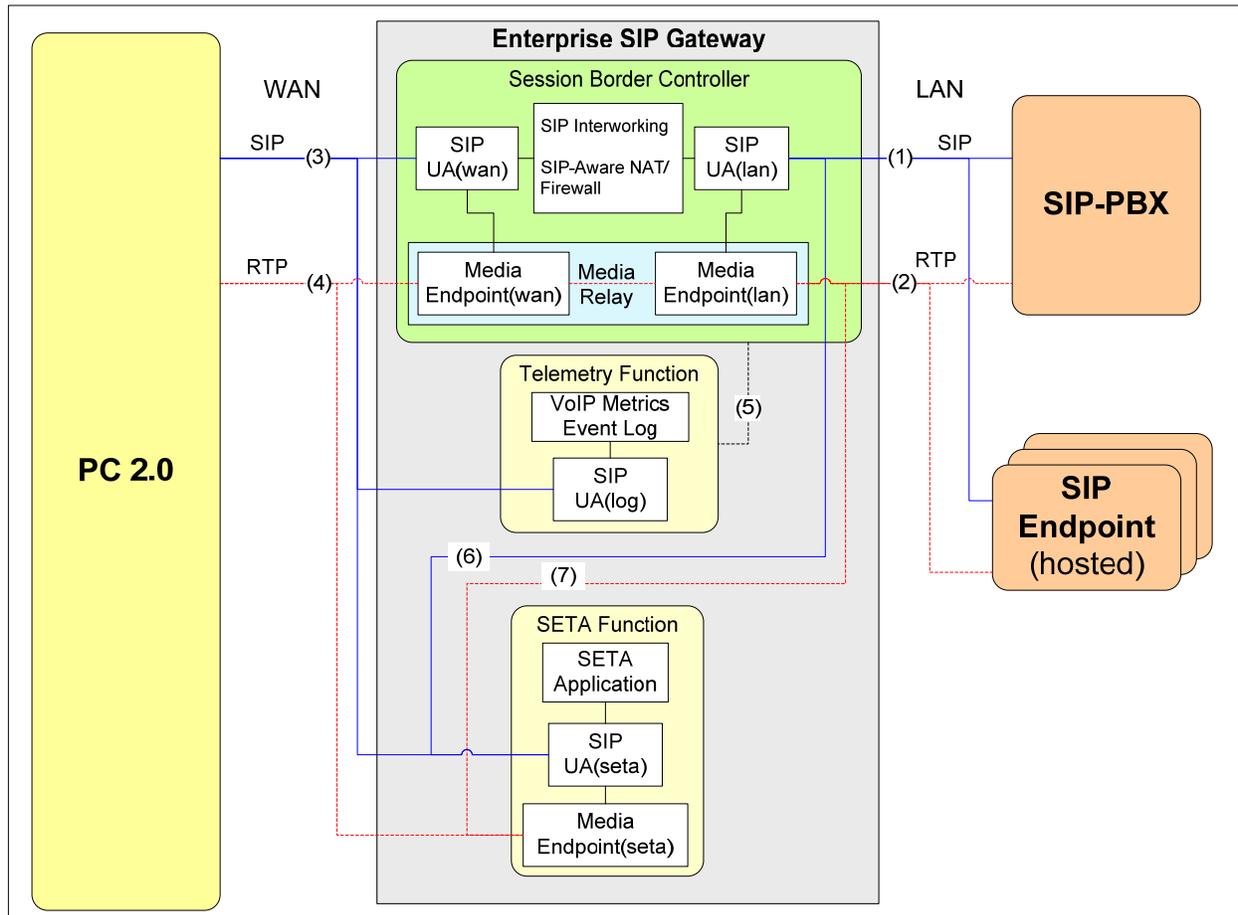


Figure 2 - Internal Architecture of ESG

The ESG SIP Endpoint Test Agent (SETA) Function and the Telemetry Function are not in the direct path of nor do they affect the business voice traffic exchanged between the Service Provider and Enterprise network. Rather, these entities serve as support functions to aid in problem detection and isolation.

6.1 Session Border Controller

The SBC supports three separate functions; SIP Interworking, Network Address Translation (NAT), and SIP firewall.

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, and SIP body (SDP) 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 Interworking 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 SIP-PBX and the Service Provider network.

As shown in Figure 2, SIP UA(wan) faces the Service Provider network, and supports SIP procedures compatible with PacketCable 2.0 on interface (3). SIP UA(lan) faces the enterprise network, and supports the SIP procedures compatible with the SIP-PBX on interface (1). These two SIP UAs are interconnected by the Interworking Function which applies SIP header manipulation rules required to achieve interworking between interfaces (1) and (3). It also enforces the SIP firewall rules.

6.1.1 SIP Interworking Function

The SIP signaling on the WAN side of the ESG must conform to PacketCable standards, as follows:

1. If the ESG is providing hosted IP Centrex service to the enterprise, then the WAN interface of the ESG must support the PacketCable Business SIP Services procedures defined in [PKT-SP-BSS].
2. If the ESG is providing SIP Trunking service to a SIP-PBX, then the WAN interface of the ESG must support the SIP-PBX requirements defined in [SIP-CONNECT1.1]. The SIPconnect 1.1 specification in turn defines two different modes of operation:
 - a) Registration Mode, where the SIP-PBX conveys its SIP signaling address to the Service Provider network using an extension of the SIP registration procedure. In this mode, the relationship between the Enterprise and the Service Provider network is a User-to-Network relationship, where the SIP-PBX access point into the PacketCable 2.0 network is via the P-CSCF.
 - b) Static Mode, where the SIP-PBX does not register with the Service Provider network. In this case, the Enterprise and Service provider network view each other as peer networks. The access point into the PacketCable 2.0 network is via the IBCF.

The SIP signaling on the LAN interface of the ESG is not defined as part of this specification effort; it must conform to whatever profile of SIP is supported by the connected SIP entity in the Enterprise network. The job of the Interworking Function is to manipulate the SIP messages and procedures, as necessary, to achieve interworking between the WAN and LAN interfaces.

The SIP Interworking Function must be scalable in terms of complexity.

In its most complex form, the SIP Interworking Function requires the SBC to act as a true B2BUA, where SIP UA(lan) terminates each dialog from the SIP-PBX and associates it with a new dialog established between SIP UA(wan) and the PacketCable 2.0 network. This case could apply for the following reasons:

1. When there are major incompatibilities between the SIP-PBX and the Service Provider network that can't be resolved with minor header manipulation (e.g., the Service Provider network expects the SIP-PBX to register, but the SIP-PBX doesn't support registration).
2. When the Service Provider decides for policy reasons to configure the SIP-PBX identity and credentials in the ESG itself. In this case, even if the SIP-PBX is compliant with the Service Provider network, the ESG registers on behalf of the SIP-PBX, generates the challenge response to Digest authentication challenges received from the network, and operates as a full B2BUA in the SIP signaling chain between the SIP-PBX and the Service Provider network.
3. Based on vendor design decision, the SBC always operates as a B2BUA.

Moving down toward the "simple" end of the complexity scale, the SIP Interworking Function manipulates SIP headers to resolve minor interworking issues, but otherwise allows the SIP-PBX to signal and establish SIP dialogs directly with the Service Provider network. In its simplest form the SIP Interworking Function is disabled; i.e., it is completely transparent to SIP signaling. This case occurs when the SIP-PBX is fully compatible with the Service Provider network, and the Service Provider decides for policy reasons to exchange SIP messages directly with the SIP-PBX. In these less-complex cases, the SBC is not required to act as a B2BUA; rather it can operate as a transparent pipe with-respect-to SIP signaling, or it can operate as a SIP Proxy as defined in [RFC 3261]. The following subsections describe these different SBC modes in more detail.

6.1.1.1 SBC Acting as B2BUA

Figure 3 expands the SBC block to show some of the data attributes contained in the internal components when the SBC is acting as a B2BUA. The network-facing SIP UA(wan) appears as PacketCable 2.0-compliant UE representing the Enterprise SIP entity (ESE).

Figure 3 shows some of the significant data attributes contained in SIP UA(wan):

- **Public Identity (IMPU):** This is the public identity of the ESE assigned by the Service Provider. If the PC 2.0 network expects the ESE to register, then this is the identity that is populated in the To header field of the REGISTER request.
- **Private Identity (IMPI):** This is the private identity of the ESE used during Digest authentication.
- **Digest password:** These are the credentials of SIP UA(wan), which are used to calculate the challenge-response during Digest authentication.
- **ESE WAN IP address:** This is the WAN-side SIP signaling IP address of SIP UA(wan) obtained from the Service Provider DHCP Server. From the PacketCable 2.0 network's perspective, this is the location of the ESE on the Service Provider network. If the PC 2.0 network expects the ESE to register, then this is the address that is populated in the Contact header field of the REGISTER request.

SIP UA(lan) faces the SIP-PBX. Figure 3 shows some of the primary data attributes owned by SIP UA(lan):

- **ESE AOR:** This is the identity of the ESE assigned by the Enterprise. It may be identical to the IMPU assigned to SIP UA(wan).
- **ESE LAN IP address:** This is the SIP signaling contact address of the SIP-PBX in the enterprise LAN. This address is either statically configured by the Enterprise on the ESG, or it is discovered by the ESG when the ESE registers.

- My IP LAN address: This is the LAN-side SIP signaling IP address of SIP UA(lan). It is either obtained from the enterprise DHCP server, or statically configured by the enterprise in the ESG.

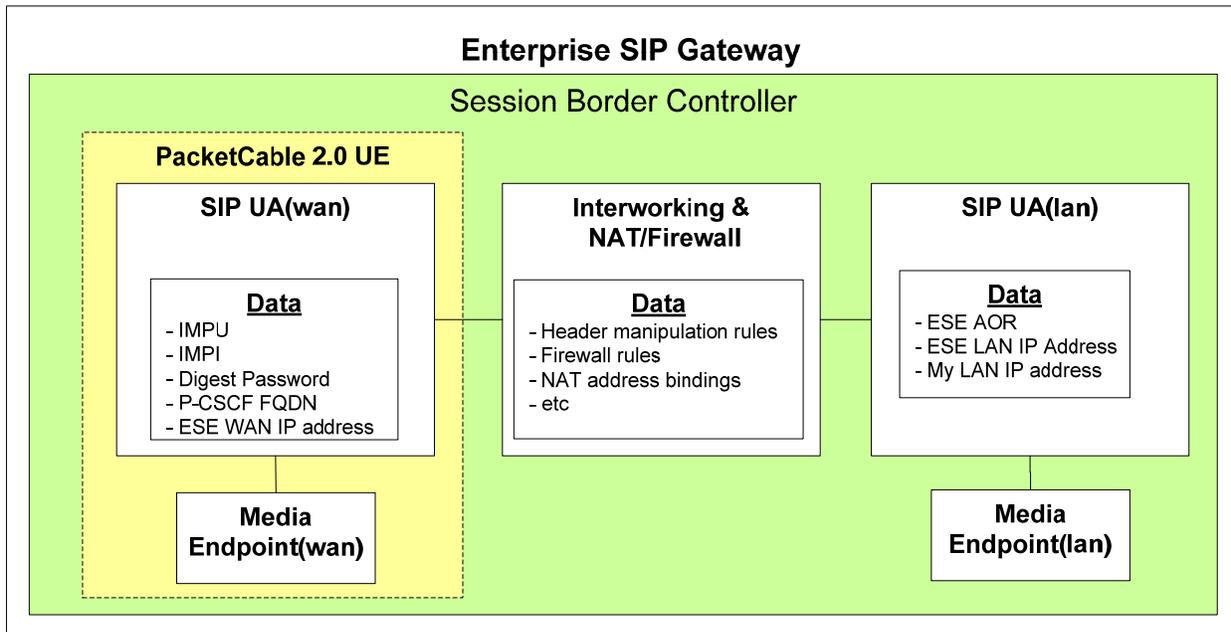


Figure 3 - Block Diagram of SBC acting as B2BUA

6.1.1.2 SBC Acting as a Transparent Pipe

When the SIP Interworking Function is configured to perform minimal or no header manipulation, and the operator policy does not require the SBC to register on behalf of the SIP-PBX, then the SBC is minimized to performing NAT/firewall functions, possibly with some minor header manipulation. As shown in Figure 4, the network-facing SIP-UA(wan) essentially moves from the SBC to the SIP-PBX, with the result that SIP dialogs are established directly between the enterprise SIP entity and a remote SIP UA within or beyond the PacketCable 2.0 network. The SBC no longer acts as a classic B2BUA; i.e., it no longer implements a back-to-back UAC/UAS that terminates and re-initiates the dialog across the demarcation point.

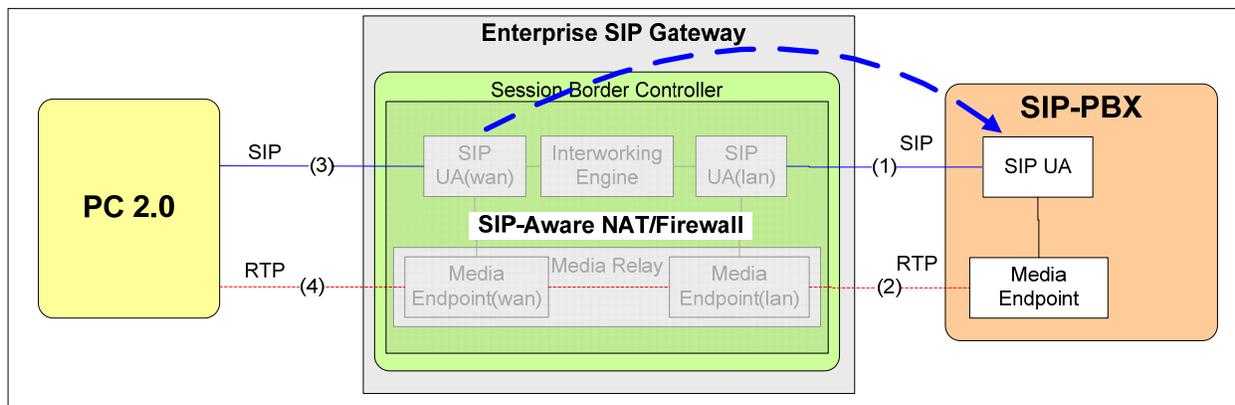


Figure 4 - SBC as NAT only

In this mode the SBC does continue to support NAT functions, and may also support minor header manipulation. But it maintains no call nor transaction state data, other than maintaining the IP address bindings required to support NAT. In fact, the box labeled "Session Border Controller" is no longer acting as an SBC as the term is normally used.

Figure 5 shows the SBC data when operating in this transparent mode. Even though the SBC does not appear as a standard SIP entity in the signaling chain, it does update SIP message data as part of its NAT function. SIP UA(wan) and SIP UA(lan) are reduced to containers of the NAT binding information. For example, the SIP UA(wan) is reduced to a simple non-SIP entity that contains sufficient information to route SIP messages to/from the ESE (i.e., an ESE WAN IP address, and possibly an ESE FQDN).

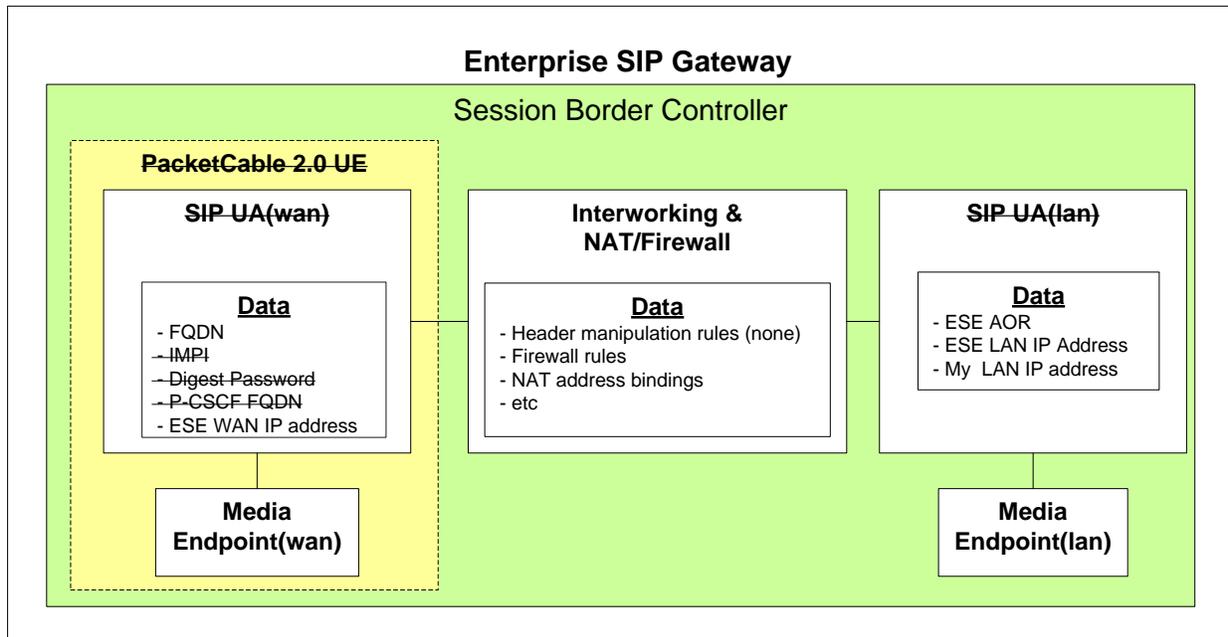


Figure 5 - Block Diagram of SBC acting as transparent pipe

6.1.1.3 SBC Acting as SIP Proxy

The SBC may also operate as a SIP Proxy. For example, the ESG can be deployed in a mode of operation where the Service Provider network views the Enterprise network as a peer network (this is referred to as the "Static mode" in [SIP-CONNECT1.1]). In this case, the ESG SBC becomes the ingress/ingress border element for the Enterprise network (i.e., an IBCF, in IMS terms). In this role the ESG could be implemented as a B2BUA that creates back-to-back SIP dialogs across the WAN/LAN boundary as described in Section 6.1.1.1. Or, if minimal SIP interworking is required as described in Section 6.1.1.2, the operator could configure the SBC to act as a SIP Proxy, where it is assigned a SIP URI that it could insert in the Record-Route or Via header fields of a dialog-initiating SIP request.

6.1.2 SIP-aware NAT/Firewall Functions

Most Enterprise networks deploy a NAT/firewall at the point where the Enterprise network connects to the Service Provider network. The NAT function translates the IP header addresses between the WAN and LAN address space, and the firewall imposes rules on access to and from the Enterprise network. Although these two functions are essential in providing secure high-speed data service to the enterprise, they can cause problems for VoIP traffic.

SIP contains IP addresses in the SIP headers and message bodies. A conventional (non-SIP-aware) NAT doesn't update these addresses, which makes it impossible to establish voice sessions between the Enterprise and the

Service Provider network. Non-SIP-aware firewalls often limit the access to the Enterprise network from the outside world, which has the undesirable side-effect of blocking new call requests from the Service Provider to the Enterprise network.

The SIP-aware NAT/firewall in the ESG solves these issues by bypassing the non-SIP-aware NAT/firewall deployed in the Enterprise for business voice traffic. The ESG SIP-aware NAT/firewall essentially takes over this role on behalf of the enterprise. The ESG NAT updates the IP header addresses just as the Enterprise NAT did. In addition, its SIP-aware capability enables it to update the IP addresses in the SIP message headers and message body, and to maintain the WAN/LAN address bindings for established calls, so that VoIP sessions can be successfully established between the Service Provider and Enterprise networks. The SIP-aware firewall understands the SIP signaling and can, therefore, open pinholes to enable incoming SIP requests from the Service Provider network to reach registered enterprise SIP entity, and to enable RTP packets to be exchanged between the Service Provider and the Enterprise for calls that have been established via SIP.

The SIP-aware NAT can also support IPv4/6 interworking, say where a Service Provider network that supports IPv6 deploys business voice service to an Enterprise network that supports IPv4.

Far end NAT traversal (where a SIP UA is on another network behind a NAT from where the ESG is located) is not in scope for the SIP NAT/firewall.

6.2 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 the ESG 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 the Figure 2, SETA can exchange its SIP messages and RTP packets directly with the Service Provider network via interfaces (3) and (4), or, as a configuration option, indirectly through the LAN-side of the SBC via interfaces (6) and (7). This later option provides a means for SETA to verify the SBC functionality. Since interfaces (6) and (7) are internal, the details of how this is accomplished are not specified.

The SETA Function must be capable of verifying connectivity with the Service Provider network for two interconnect relationships between the Service Provider and Enterprise network; for the user-to-network case where the access is via the P-CSCF, and the peer-to-peer case where the access is via the IBCF (see Section 6.1.1. for more information on the Enterprise-SP network interconnect modes). In the user-to-network case, the SETA appears as a PacketCable 2.0-compliant UE that registers with the core network via the P-CSCF. In the peer-to-peer case, the SETA appears as a remote endpoint in a peer network, and therefore, does not register with the core Service Provider network.

6.2.1 Base Functionality

The SETA line has much in common with a PacketCable 2.0 UE. It has a public identity, and can originate and receive calls using the base PC 2.0 UE procedures. The unique capability of the SETA is that the traditional analog loop interface that is used to originate and terminate calls is replaced with management interface.

6.2.2 RTP Loopback

SETA should support RTP Loopback which is defined in the Loopback Test Capacity section of [PKT-RST-E-DVA]. The ESG test line only needs to support the "rtp-pkt-loopback" mode, where it acts as a packet reflector (i.e., no decode/encode). This allows the network operator to test the voice path between the Service Provider network and the ESG, and to collect VoIP performance statistics for voice connections on the Service Provider side of the demarcation point.

6.2.3 Test Call Origination and Termination

The SETA provides management controls that enable the Service Provider to configure a dial string, call duration in seconds, and a wave file which is to be transmitted over the test call session. The operator can instruct the SETA to initiate the test call on-demand or at periodic intervals. When so instructed, the SETA establishes the call with the remote endpoint, collects call statics (error events, VoIP metrics), and terminates the call once the call duration has expired.

The SETA reports the call disposition of test calls (success, call setup time, post dial delay, busy, blocked, dropped, etc.) and call quality (packet loss, latency, jitter, etc.) to the Service Provider through the management interface.

6.3 Telemetry Function

6.3.1 Voice Statistics

6.3.1.1 Collecting VoIP Metrics

Since the ESG acts as a media relay element between the enterprise SIP entity (SIP-PBX or hosted SIP endpoint) and the remote SIP endpoint, it sees all business voice RTP and RTCP traffic between the Enterprise and the Service Provider network. The ESG is, therefore, in a position to gather the following statistics related to network performance. The ESG is not an actual media endpoint that decodes the incoming RTP stream, and therefore, it cannot accurately generate all the VoIP metrics parameters defined in [RFC 3611]. However, it can generate useful "pseudo metrics" that provide a good indication of voice quality.

The ESG will collect the following metrics:

- Packet Interarrival Jitter calculated as 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. Packet Interarrival Jitter is calculated 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.
- Packet Loss. Two packet loss estimates must be calculated. The first is the fraction of RTP packets lost during a configurable time interval T_{pl} , which 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 T_{pl} is started at the beginning of RTP transmission and the count of the fraction of RTP packets lost is reset every T_{pl} seconds. The second estimate of packet loss is 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. The fraction of RTP packets lost and the cumulative number of RTP packets lost must be calculated 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.
- Round-Trip Propagation Delay. The ESG measures two Round-Trip Delays. Both delays rely on the sending of RTCP reports by the SIP endpoints and both delays must be calculated if RTCP reports are available. One delay relates to the leg from the ESG to the SIP-PBX endpoint and back to the ESG, while the other relates to the leg from the ESG to the remote SIP endpoint and back to the ESG. Both these delays are calculated as in [RFC 3550]. Specifically, the delay is based on the timing of the sending and reception of RTCP reports at the ESG and the DLSR (Delay Since Last SR) parameter received in RTCP reports from the SIP endpoints. 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 report 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 consider implementing some error checking to ensure that patently erroneous values of delay are not reported.

The ESG will collect the above metrics for two RTP streams per call; the incoming RTP stream from the local enterprise endpoint, and the RTP stream from the remote 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-XR reports are available, 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.

6.3.1.2 Reporting VoIP Metrics

The ESG can provide the collected voice metrics to the Service Provider via the following mechanisms:

- Local log accessible from a Web GUI
- Sending call statistics to an external Syslog server
- SIP PUBLISH

These mechanisms must be 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.

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.

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 [ID-SIP-RTCP]. The back-office server Collector Function that receives the VoIP metrics reports is referred to as the '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 [ID-SIP-RTCP]. It is informed of the contact address of the collector as part of the registration process. The ESG must populate the request URI in the PUBLISH request with the collector address. Separate PUBLISH messages will be sent 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. It will be possible to separate which direction each PUBLISH message relates to by examining the 'RemoteID' parameter within the PUBLISH message body.

There are three types of metrics reports using 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 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 collector on a regular basis for all active sessions. Alert reports will also be sent to the 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 collector. As noted in [ID-SIP-RTCP], 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 collector. The debug mode should, therefore, be employed as a temporary means of troubleshooting rather than a normal mode of operation.

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 collector. These will include:

- NetworkPacketLossRate (NLR)

- BurstLossDensity (BLD)
- BurstDuration (BD)
- GapLossDensity (GLD)
- GapDuration (GD)
- MinimumGapThreshold (GMIN)
- RoundTripDelay (RTD)
- InterarrivalJitter (IAJ)

These parameters will be calculated and sent in separate PUBLISH messages for both the upstream and downstream legs of the call. It is also possible to populate the 'RemoteMetrics' block of the PUBLISH message from RTCP-XR reports that are sent by the remote SIP endpoint or the SIP-PBX endpoint if these RTCP-XR reports are available.

6.3.2 SIP and RTP Trace

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. Each of these traces includes all SIP messages up to the current time for the corresponding call. The SIP trace function must show the SIP procedures on both the WAN and LAN side of the SBC, so that the Service Provider can verify the header manipulation and interworking procedures applied by the Interworking Function.

The ESG must provide Service Provider management controls to instruct the ESG 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. 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 can be stored in the data format defined for a popular capturing tool such as WireShark. They can be queried either locally or remotely by the operator.

6.4 Quality of Service

QoS in the DOCSIS access network is provided by PacketCable Multimedia [PCMM] and the underlying DOCSIS layer two infrastructure [MULPI3.0].

DOCSIS has the concept of 'Service Flows' and 'Classifiers'. A Classifier is composed of the fields in an IP header: <From-IP>, <From-Port>, <To-IP>, <To-Port>, and <DiffServeCodePoint>. Any of these Classifier fields can be wild-carded. One or more Classifiers can point to a Service Flow where Service Flow is a specific upstream or downstream QoS or filtering mechanism defined in the DOCSIS standard [MULPI3.0].

PacketCable Multimedia has the notion of a 'Gate'. Gates are unidirectional. To set up a two-way flow, an application needs to use PCMM signaling to install two separate gates for the upstream and the downstream. The information in the PacketCable Multimedia signaling maps directly onto DOCSIS classifiers and services.

In the downstream direction, the CMTS provides a priority queuing mechanism based on DiffServ where SIP, RTP, and RTCP traffic are given priority over best-effort traffic. SIP signaling, RTP, and RTCP can be given separate TOS/DiffServ Code Points (DSCP) to allow the CMTS to place RTP at a higher queuing priority than SIP or RTCP. These DSCP are typically not allowed to be used by anything else in the network. In practice the simplest approach for voice applications is to statically configure the downstream so the DSCP used for voice and voice signaling are given high priority.

In the upstream direction, DOCSIS provides a mechanism called Unsolicited Grant Service (UGS). In UGS, the CMTS issues a grant to the cable modem at a fixed time interval to transmit a fixed amount of upstream data. Typically, the interval is 20 milliseconds and the amount of data is the size of an entire RTP frame including MAC, IP, UDP, and RTP headers. Using PCMM, the application installs a gate for a specific RTP flow where the service is UGS and the classifier is the <To-IP>, <To-Port>, and <DSCP> of the distant end media gateway with <From-IP> and <From-Port> wildcarded.

DOCSIS cable modem implementations have finite resources. As Business Voice traffic increases, these resource limits can limit the number of simultaneous calls supported. In particular, there is typically a limit on the number of service flows supported by a cable modem and the number of classifiers supported by a cable modem. There are typically many more service flows than classifiers. To work around this limitation, DOCSIS has the notion of "Multiple Grants per Interval" (MGPI) in Unsolicited Grant Service (UGS). For example, in a single 20 millisecond interval in a service flow, the CMTS can be told to issue 10 grants for different RTP streams. Each RTP stream would have a different classifier which points at the same service flow.

Given this underlying mechanism, the critical function of the ESG is to ensure that the SIP, RTCP, and RTP packets on the upstream are all using the correct DiffServ Code Points. For RTP traffic, it is also critical that the full upstream IP header match what was signaled in the SDP since the application could be using all of <From-IP>, <From-Port>, <To-IP>, <To-Port>, and <DiffServeCodePoint> as the classifier.

6.5 Security Services

A full set of relevant threats to the PC 2.0 network has been studied, classified and documented in [PKT-TR-SEC]. In general, all security considerations outlined in [PKT-TR-SEC] are applicable to the ESG with respect to communications to the MSO network.

6.5.1 SIP Signaling Security

The ESG security requirements and procedures differ depending on which SIP Signaling interface is considered.

Security service for the SIP interface between the ESG and the Service Provider network are provided with the principal goal of protecting the Service Provider's SIP Core network from unauthorized use, and from threats such as Denial of Service (DoS) attacks. The primary security mechanisms used on this interface are SIP Digest authentication, and Transport Layer Security (TLS). The ESG must support the SIP-PBX security requirements in [ID-SIP-RTCP] when providing SIP Trunking service, and must support the E-DVA security requirements specified in [PKT-RST-E-DVA], when providing IP hosted Centrex service.

Security service requirements for the SIP interface between the ESG and the enterprise SIP endpoints are outside the scope of this specification effort. However, to protect the MSO network to the fullest extent possible, it is highly recommended that this interface is protected by security services such as SIP Digest and TLS.

6.5.2 Security Services for RTP Media

Security Services for RTP/RTCP are outside of the scope of this specification effort.

6.5.3 Security Services for Provisioning & Management Interfaces

As an implementation option, the Security Services for the Provisioning and Management interfaces between the ESG and the Service Provider network may align with the security procedures specified for PacketCable 2.0 in [PKT-RST-E-DVA] and [PKT-RST-EUE-PROV].

6.6 Operations, Administration, Management, and Provisioning

This section contains a general description of the mechanisms and operations used for Provisioning of the ESG. Such operations cover the Configuration, Management, and Event Reporting functions, required by the Service Provider. All these functions taken in their entirety are referred to in this section as "Provisioning".

As shown in Figure 2, the ESG contains a SIP UA(wan), SIP UA(seta), and SIP UA(telemetry) that have much in common with the SIP UA defined for the E-DVA. Therefore, as an implementation option, the ESG can use the provisioning framework defined for the PacketCable 2.0 E-DVA in [PKT-EUE-PROV], [PKT-RST-EUE-PROV], and [PKT-RST-E-DVA] specifications. These procedures could be applied to both the embedded ESG and to the stand-alone ESG since the stand-alone device is never deployed behind a NAT (i.e., the stand-alone ESG always obtains a publically reachable WAN IP address from the Service Provider network). This SNMP/MIB-based implementation option would cater to those Service Providers that want to extend their E-DVA provisioning processes to support the ESG.

As an implementation option, the ESG can support a different provisioning framework (e.g., XML-based), that could be deployed by those Service Providers that don't plan to extend their E-DVA provisioning processed to support the ESG.

The PacketCable 2.0 provisioning specifications [PKT-UE-DATA] and [PKT-EUE-DATA] define a Data Model and MIBs for E-DVA provisioning. This ESG specification effort will not extend these E-DVA provisioning specifications to support the new data attributes required by the ESG. Rather, the ESG specifications will define a generic data model that is independent of the underlying provisioning framework. This general data model can then be used to generate the specific data format and syntax required by whatever provisioning framework is supported by the ESG.

The ESG specification effort covers the provisioning requirements for the ESG only. The provisioning requirements for SIP-PBX are out of scope of this document. This document makes no assumptions as to what provisioning procedures and data objects are used to enable the functionality of the SIP-PBX.

Appendix I Acknowledgements

We wish to thank the following 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):

Peter Som De Cerff, Adtran

Eugene Nechamkin, Broadcom

Gordon Li, Broadcom

Doug Wadkins, Edgewater

Guhan Parthasarathy, Global Edge Software Ltd

Lee Valerius, Huawei

Harprit Chhatwal, InnoMedia

Geoff Devine, SMC Networks

Satish Kumar, TI

The following CableLabs staff is thanked for their direct contributions to this Technical Report: Eduardo Cardona and Vikas Sarawat.

David Hancock and the Business Services Team (CableLabs)
