

Dual-Stack IPv6 Architecture Technical Report

PKT-TR-DS-IP6-C01-140314

CLOSED

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. 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., 2011-2014

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 affiliated 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.

Abstract

This technical report describes the IPv6 Dual-Stack architecture in PacketCable™ networks, including the major system components, the various functional groupings and the network interfaces necessary for IPv6 Dual-Stack support of end-to-end communications. A number of potential scenarios for adopting or migrating to IPv6 Dual-Stack are presented along with a high-level theory of operation for IPv6 Dual-Stack User Equipment. This technical report also initiates an inventory of the PacketCable specifications that may be impacted by IPv6 Dual-Stack.

The intended audience for this document includes developers of equipment intended to be conformant to PacketCable specifications, and network architects who need to understand the overall PacketCable architecture framework.

This technical report contains the following information:

- Description of IPv6 Dual-Stack support in PacketCable networks;
- Impact of IPv6 Dual-Stack for all PacketCable networks element, with a special focus on the client side (User Equipment);
- IPv4-IPv6 scenarios for PacketCable Dual-Stack deployments over DOCSIS cable access networks;
- Theory of operations for IPv6 Dual-Stack User Equipment clients including a high-level description of PacketCable provisioning of IPv6/IPv4 Dual-Stacked User Equipment;
- Preliminary analysis of the IPv6 Dual-Stack impact on PacketCable specifications.

The PacketCable specifications take precedence over this technical report if the technical report contradicts any specification requirements.

Document Status Sheet

Document Control Number: PKT-TR-DS-IP6-C01-140314

Document Title: Dual-Stack IPv6 Architecture Technical Report

Revision History: V01 – 08/25/11
C01 – 03/14/14

Date: March 14, 2014

Trademarks

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

Contents

1	INTRODUCTION	1
1.1	PacketCable Overview	1
1.2	PacketCable IPv6 Overview	1
1.3	IPv6 Dual-Stack	1
1.4	Document Organization	2
2	REFERENCES	3
2.1	Informative References	3
2.2	Reference Acquisition	4
3	TERMS AND DEFINITIONS	5
4	ABBREVIATIONS AND ACRONYMS	7
5	FUNCTIONAL REQUIREMENTS	8
5.1	UE Provisioning Requirements	8
5.1.1	<i>Review of IPv6-only eUE Provisioning</i>	<i>8</i>
5.1.2	<i>eUE Dual-Stack (IPv4 & IPv6) Provisioning Requirements</i>	<i>8</i>
5.1.3	<i>eUE SIP Registration</i>	<i>9</i>
5.2	Session Signaling and Media Requirements	9
5.2.1	<i>Dual-Stack Interworking</i>	<i>9</i>
6	PROVISIONING	10
6.1	Dual Stack Provisioning requirements	10
6.2	SIP Registration configuration	10
6.3	Preferred Media Advertisement	10
7	SESSION SIGNALING AND MEDIA	11
7.1	IP Version Interworking at the IP Layer	11
7.1.1	<i>IP Layer in SIP Signaling Plane – within PacketCable 2.0 Core</i>	<i>11</i>
7.1.2	<i>IP Layer in Signaling Plane – between Endpoints and PacketCable 2.0 Core</i>	<i>11</i>
7.1.3	<i>IP Layer in Media Plane</i>	<i>11</i>
7.2	IP Version Interworking at the Signaling Layer	12
7.2.1	<i>Signaling Layer – SIP Header Fields</i>	<i>12</i>
7.2.2	<i>Signaling Layer – SDP Body</i>	<i>12</i>
7.3	IP Version Interworking Procedures in PacketCable 2.0	12
7.3.1	<i>IP Version Interworking in a Single-Stack Environment</i>	<i>12</i>
7.3.2	<i>Negotiating Media IP Version among Dual-Stack Endpoints</i>	<i>13</i>
7.3.3	<i>IPv4/6 Interworking in a Mixed Single/Dual-Stack Environment</i>	<i>18</i>
7.4	Quality of Service and PCMM	21
8	USE CASES	22
8.1	Ideal IPv4-to-IPv6 Migration Strategy	22
8.1.1	<i>Migrate the IP Core Network to Dual-Stack</i>	<i>22</i>
8.1.2	<i>Migrate the PacketCable 2.0 Core Network to Dual-Stack</i>	<i>22</i>
8.1.3	<i>Migrate the PacketCable 2.0 Endpoints to Dual-Stack</i>	<i>23</i>
8.1.4	<i>Deploy Single-Stack IPv6-only Endpoints</i>	<i>23</i>
8.2	Deviations from the Ideal IPv4-to-IPv6 Migration Path	24
8.3	Call Flow Diagrams	24
8.3.1	<i>Establishing Calls over a Dual-Stack PacketCable 2.0 Network</i>	<i>25</i>
8.3.2	<i>Establishing Calls over a Mix of Single- and Dual-Stack Networks</i>	<i>32</i>
8.3.3	<i>Interworking Between PacketCable 1.5 and 2.0</i>	<i>35</i>
8.3.4	<i>Two Dual-Stack Networks Connected by an IPv4 Transit Network</i>	<i>36</i>

APPENDIX I ACKNOWLEDGEMENTS	37
---	-----------

Figures

Figure 1 - eUE Provisioning Sequence Diagram	8
Figure 2 - Using ICE-lite to Negotiate Media IP-version	15
Figure 4 – ALG/Media-Relay Support of ICE-lite	20
Figure 5 - Dual-Stack IP Network	22
Figure 6 - Dual-Stack PacketCable 2.0 Network	23
Figure 7 - Dual-Stack Network with Mix of Dual-Stack and IPv4-only Endpoints	23
Figure 8 - Dual-Stack Network with Mix of Dual-Stack and IPv6-only Endpoints	24
Figure 9 - Message Flow Template	25
Figure 10 - End-to-End Dual-Stack Network Supporting Mix of Single and Dual-Stack UEs	25
Figure 11 - IP Version Interworking Between Single-Stack UEs in a Dual-Stack Network	26
Figure 12 - Single-Stack UE-1 Calls Single-Stack UE-2	27
Figure 13 - Dual-Stack UE-1 Registered using IPv4 Calls Dual-Stack UE-2 Registered using IPv6	29
Figure 14 - Dual-Stack UE Registered Using IPv6 Calls Single-Stack IPv4 UE	30
Figure 15 - Dual-Stack UE Registered Using IPv4 Calls Single-Stack IPv6 UE	31
Figure 16 - Single-Stack IPv6 UE Calls Single-Stack IPv4 UE	32
Figure 17 - Mixed Single and Dual-Stack Networks	33
Figure 18 - Dual-Stack UE Calls Single-Stack IPv4 UE over Mixed IPv4/6 Network	34
Figure 19 - Single-Stack IPv6 UE Calls Single-Stack IPv4 UE over Mixed Network	35
Figure 20 - Interworking Between PacketCable 1.5 and Dual-Stack PacketCable 2.0 Networks	36
Figure 21 - Dual-Stack Networks Interconnected by Single-Stack Transit Network	36

Tables

Table 1 - Dual-Stack Media Interworking Scenarios	9
---	---

1 INTRODUCTION

This technical report provides an overview of the IPv6 Dual-Stack support in PacketCable. Specifically, it describes the way in which PacketCable networks may adopt or migrate to IPv6 using a phased approach utilizing dual-stack functionality to provide end-to-end IP connectivity support to PacketCable applications.

To aid the reader in understanding the PacketCable IPv6 Dual-Stack architecture, the impact of IPv6 Dual-Stack support in the PacketCable Reference Architecture functional components is discussed in this document along with the affected protocol interfaces. CableLabs has issued this technical report and associated specifications to facilitate design and field-testing leading to the manufacture and interoperability of conforming hardware and software by multiple vendors.

1.1 PacketCable Overview

The PacketCable Architecture Framework Technical Report [ARCH-FRM TR] describes a set of functional groups and logical network entities grouped by service functions, as well as a set of reference points that support the information flows exchanged between these functional groups and network entities.

Except where explicitly qualified as PacketCable 1.0, or PacketCable 1.5, the term PacketCable generally refers to PacketCable 2.0.

1.2 PacketCable IPv6 Overview

The Internet Protocol, version 6 is defined by IETF; a partial list of the IETF IPv6 specifications includes [RFC 2460], [RFC 4861], [RFC 4862], [RFC 4443], [RFC 3513], and [RFC 4213]. The CableLabs DOCSIS® 3.0 specifications add full support for IPv6 on cable access networks [MULPI].

The support of IPv6 in PacketCable networks impacts all functional groupings, from the User Equipment (UE) and the Local Network the UE uses to connect to the access network, to the Access Network, Edge and Operational Support areas, including critical areas like the Core functional grouping (3GPP SIP-based IMS network elements), PacketCable Multimedia and the Interconnect area allowing the interconnection of networks using different versions of IP. A presentation of the IPv6 support in the functional groupings defined in [ARCH-FRM TR] is provided in Section 5.2.

IPv6 also affects many protocol interfaces defined in PacketCable, including the provisioning, activation, configuration and management of an IPv6 UE, the signaling and QoS interface specifications for UE registration and communication establishment with the IMS Edge and Core components, the media stream transport over IPv6 through NATs and Firewall devices, etc.

It is assumed that UEs implement both IPv4 and IPv6 and both stacks can be configured to operate on the device. An IPv6/IPv4 UE initiates PacketCable provisioning for IPv6 and IPv4 to acquire IP address information, discover its P-CSCF and retrieve configuration parameters. The selection of the IP version a UE should use is based on a combination of its local network capabilities, the DNS information returned for the P-CSCF, the use of [RFC 3484] to make source and destination address selection and the provided configuration data.

1.3 IPv6 Dual-Stack

This document describes the addition of dual-stack operation where the UE acquires both an IPv4 and IPv6 address during provisioning and uses both IPv4 and IPv6 for media as needed. This facilitates an easier transition between legacy IPv4-only networks and future IPv6-only networks. While dual-stack operation does not inherently save any IPv4 address resources, it does allow for the deployment of IPv6-only devices in the future. Therefore dual-stack operation is a bridge between the current IPv4-only world and the future IPv6-only world, allowing backwards compatibility and a reduced dependence on Translational Gateways.

NOTE: Within the context of this document, a UE is a PacketCable 2.0 SIP endpoint that is implemented by an E-DVA or at the WAN interface of an Enterprise SIP Gateway (ESG).

1.4 Document Organization

The purpose of this technical report is to describe the overall end-to-end technical solution for dual-stack; how the PacketCable network uses the dual-stack capability to facilitate migration to IPv6.

Chapter 5 describes the high-level functional requirements that serve as input into the dual-stack project. These functional requirements must be supported by the procedures and mechanisms described in the remainder of the document.

Chapter 6 describes the new provisioning procedures that must be supported by a dual-stack embedded UE (eUE). The associated normative provisioning requirements for the eUE are defined in the following two PacketCable specifications;

- [PKT-EUE-PROV] specifies the changes to the eUE provisioning process for dual-stack.
- [PKT-EUE-DATA] defines eUE data attributes for control of dual-stack operation and event messages for the added provisioning processes.

Since the eUE is the base entity for all embedded SIP endpoints defined in PacketCable 2.0, these new normative provisioning requirements apply to all embedded PacketCable 2.0 endpoints; namely, the E-DVA and the two embedded versions of the ESG – the ESG E-DVA and the E-ESG. Since it is not an eUE device, the stand-alone ESG (the S-ESG) dual-stack provisioning requirements are specified separately, in [PKT-ESG].

Chapter 7 describes the new session signaling procedures that are used at the SIP layer by a dual-stack UE to determine its SIP signaling address and to negotiate the media IP version during session establishment. The associated normative session signaling requirements are defined in three PacketCable specifications:

- [PKT 24.229] specifies the base UE normative requirements for negotiating the media IP version when the UE supports dual-stack. These requirements are defined at an "optional to support" level (i.e., not mandatory) so that they can be introduced into 3GPP IMS in the future.
- [PKT-RST-E-DVA] specifies the normative requirements for the E-DVA to determine its SIP signaling address and to negotiate the media IP version during session establishment. This specification raises the level of [PKT 24.229] dual-stack UE requirements from "optional" to "mandatory" for the E-DVA.
- [PKT-ESG] specifies the normative requirements for the ESG to determine its WAN SIP signaling address and to negotiate the WAN media IP version during session establishment. This specification raises the level of [PKT 24.229] dual-stack UE requirements from "optional" to "mandatory" for the ESG.

Chapter 7 also describes how the new dual-stack procedures defined above complement the existing IP version interworking mechanisms supported by PacketCable 2.0.

Chapter 8 illustrates how the procedures and mechanisms described in Chapter 7 are used to support a number of different use-case scenarios.

2 REFERENCES

2.1 Informative References

This Technical Report uses the following informative references:

- [ARCH-1.5 TR] PacketCable 1.5 Architecture Framework Technical Report, PKT-TR-ARCH1.5-V02-070412, April 12, 2007, Cable Television Laboratories, Inc.
- [ARCH-FRM TR] PacketCable Architecture Framework Technical Report, PKT-TR-ARCH-FRM-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
- [MM ARCH] Multimedia Architecture Framework, PacketCable Technical Report, PKT-TR-MM-ARCH-V03-091029, October 29, 2009, Cable Television Laboratories, Inc.
- [MM WS] PacketCable Multimedia Web Service Interface Specification, PKT-SP-MM-WS-I03-091029, October 29, 2009, Cable Television Laboratories, Inc.
- [MULPI] Data-Over-Cable Service Interface Specifications DOCSIS 3.0, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-I23-131120, November 20, 2013, Cable Television Laboratories, Inc.
- [PCMM] PacketCable Multimedia Specification, PKT-SP-MM-I06-110629, June 29, 2011, Cable Television Laboratories, Inc.
- [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.
- [PKT-CPD] PacketCable Control Point Discovery Specification, PKT-SP-CPD-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
- [PKT-ESG] PacketCable Enterprise SIP Gateway Specification, PKT-SP-ESG-I01-101103, November 3, 2010, 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.
- [RFC 2460] IETF RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, December 1998.
- [RFC 2529] IETF RFC 2529, Transmission of IPv6 over IPv4 Domains without Explicit Tunnels, March 1999.
- [RFC 2661] IETF RFC 2661/BCP0049, Layer Two Tunneling Protocol L2TP, August 1999.
- [RFC 2874] IETF RFC 2874, DNS Extensions to Support IPv6 Address Aggregation and Renumbering, July 2000.
- [RFC 3056] IETF RFC 3056, Connection of IPv6 Domains via IPv4 Clouds, February 2001.
- [RFC 3086] IETF RFC 3086, Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification, April 2001.
- [RFC 3261] IETF RFC 3261 SIP: Session Initiation Protocol, June 2002.
- [RFC 3315] IETF RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), July 2003.
- [RFC 3316] IETF RFC 3316, Internet Protocol Version 6 (IPv6) for Some Second and Third Generation Cellular Hosts, April 2003.

- [RFC 3484] IETF RFC 3484, Default Address Selection for Internet Protocol version 6 (IPv6), February 2003.
- [RFC 3513] IETF RFC 3513, Internet Protocol Version 6 (IPv6) Addressing Architecture, April 2003.
- [RFC 3551] IETF RFC 3551/STD0065, RTP Profile for Audio and Video Conferences with Minimal Control, July 2003.
- [RFC 3596] IETF RFC 3596, DNS Extensions to Support IP Version 6, October 2003.
- [RFC 3596] IETF RFC 3596, DNS Extensions to Support IP Version 6, October 2003.
- [RFC 3646] IETF RFC 3646, DNS Configuration Options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6), December 2003.
- [RFC 3736] IETF RFC 3736, Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6, April 2004.
- [RFC 3931] IETF RFC 3931 Layer Two Tunneling Protocol - Version 3 (L2TPv3), March 2005.
- [RFC 4075] IETF RFC 4075, Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6, May 2005.
- [RFC 4213] IETF RFC 4213, Basic Transition Mechanisms for IPv6 Hosts and Routers, October 2005.
- [RFC 4215] IETF RFC 4215, Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks, October 2005.
- [RFC 4443] IETF RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification, March 2006.
- [RFC 4472] IETF RFC 4472, Operational Considerations and Issues with IPv6 DNS, April 2006.
- [RFC 4779] IETF RFC 4779, ISP IPv6 Deployment Scenarios in Broadband Access Networks, January 2007.
- [RFC 4861] IETF RFC 4861, Neighbor Discovery for IP version 6 (IPv6), September 2007.
- [RFC 4862] IETF RFC 4862, IPv6 Stateless Address Autoconfiguration, September 2007.
- [RFC 5245] 5 IETF RFC 5245 Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer, April 2010.
- [RFC 5905] IETF RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification, June 2010.
- [RFI2.0] Data-Over-Data Service Interface Specifications DOCSIS 2.0, Radio Frequency Interface Specification, CM-SP-RFiv2.0-C02-090422, CableLabs, Inc.
- [TS 23.221] 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Architectural requirements (Release 6), V6.3.0, June 2004.
- [TS 23.228] 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; IP Multimedia Subsystem (IMS); Stage 2 (Release 7), June 2005.

2.2 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone 303-661-9100; Fax 303-661-9199; Internet: <http://www.packetcable.com/>; <http://www.cablemodem.com/>.
- 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, Internet, <http://www.ietf.org/>.
- 3rd Generation Partnership Project (3GPP), ETSI Mobile Competence Centre, 650 route des Lucioles, 06921 Sophia-Antipolis Cedex, France, Internet: <http://www.3gpp.org/>.

3 TERMS AND DEFINITIONS

PacketCable Specifications and Technical Reports use the following terms and definitions:

Address family	The IP version of an IP address (i.e., IPv4 or IPv6).
Authorized Identity	An instance of an 'Authorized Identity' in a PacketCable network is a representation of an allowed pairing between a Private Identity and a Public Identity.
Client	A client is an application that runs on a device, interacts with the network and provides interfaces to users or entities.
Contact Address	The URI of a User Agent on the network. Contact addresses, in the context of PacketCable are often, but not always, addresses used to deliver requests to a specific User Agent.
Device	A physical piece or assemblage of equipment capable of hosting one or more PacketCable Clients.
DHCPv6	IPv6 version of the Dynamic Host Configuration Protocol.
Dual-Stack	See Dual-Stack mode.
Dual-Stack mode	A configuration option where a device acquires and uses IP addresses from both address families.
Dual-stack node	An IPv6/IPv4 host or router that operates with both the IPv6 and IPv4 stacks enabled.
E.164	E.164 is an ITU-T Recommendation which defines the international public telecommunication numbering plan used in the PSTN and other data networks.
Identity Credentials	A collection of the information needed to perform authentication of a private identity. The actual information depends on the authentication mechanism.
IPv4	Internet Protocol version 4.
IPv4 node	A host or router that implements IPv4. IPv6/IPv4 and IPv4-only nodes are both IPv4 nodes.
IPv4-only node	A host or router that implements only IPv4. An IPv4-only node does not understand IPv6, e.g., an IPv4-only UE that implements only IPv4. This definition is consistent with [RFC 4213].
IPv6	Internet Protocol version 6.
IPv6 node	A host or router that implements IPv6. IPv6/IPv4 and IPv6-only nodes are both IPv6 nodes.
IPv6/IPv4 node	A host or router that implements both IPv4 and IPv6. For e.g., an IPv6/IPv4 UE implements both IPv4 and IPv6. This definition is consistent with [RFC 4213].
IPv6-only node	A host or router that implements IPv6 and does not implement IPv4. For e.g., an IPv6-only UE implements IPv6 and does not implement IPv4. This definition is consistent with [RFC 4213].
Multimedia session	A set of multimedia senders and receivers and the data streams flowing from senders to receivers. A multimedia conference is an example of a multimedia session.
PacketCable Administrative Domain	An administrative domain is a collection of PacketCable elements, as described by the PacketCable Subscriber Model, managed by a single administrative authority.
PacketCable Service Provider	A cable operator operating one or more independent PacketCable Administrative Domains.

PacketCable Service Provider DNS Domain	A DNS Domain Name that is owned and managed by a PacketCable Administrative Domain. It is used to form SIP URIs that convey Public Identifiers.
Private Identity	A logical identity for purposes of authentication and authorization of a User.
Proxy Server	An intermediary entity that acts as both a server and a client for the purpose of making requests on behalf of other clients. A proxy server primarily plays the role of routing, which means its job is to ensure that a request is sent to another entity "closer" to the targeted user. Proxies are also useful for enforcing policy (for example, making sure a user is allowed to make a call). A proxy interprets, and, if necessary, rewrites specific parts of a request message before forwarding it.
Public Identifier	An identifier used to reference a Public Identity.
Public Identity	A logical identity for purposes of communication with a User. A Public Identity cannot be directly referenced - it must be referenced by using one of its Public Identifiers.
Server	A network element that receives requests in order to service them and sends back responses to those requests. Examples of servers are proxies, User Agent servers, redirect servers, and registrars.
SIP User Agent	Same as 'User Agent'.
Subscriber	The primary billed entity for a Subscription.
Subscription	A contract between a 'Subscriber' and a 'PacketCable Administrative Domain'.
User	A person who, in the context of this document, uses a defined service or invokes a feature on a Device.
User Agent (UA)	A software entity contained in a device that acts on behalf of the user to send requests to and receive responses from the network for a particular application. In the context of this document, a UA refers to a SIP User Agent as defined by [RFC 3261].

4 ABBREVIATIONS AND ACRONYMS

This Technical Reports use the following abbreviations and acronyms:

ALG	Application Level Gateway
CSCF	Call Session Control Function
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DOCSIS®	Data-Over-Cable Service Interface Specifications
eCM	embedded Cable Modem
E-DVA	Embedded Digital Voice Adaptor
E-ESG	Embedded ESG
ESG	Enterprise SIP Gateway
eUE	embedded UE
E-UE	Embedded User Equipment
ICE	Interactive Connectivity Establishment protocol as defined in [RFC 5245]
ICE-Lite	The “LITE” implementation of ICE
P-CSCF	Proxy CSCF
SAAC	Stateless Address Auto Configuration
S-CSCF	Serving CSCF
SNMP	Simple Network Management Protocol
TFTP	Trivial File Transfer Protocol
ToD	Time of Day
TrGW	Transition Gateway
UE	User Equipment

5 FUNCTIONAL REQUIREMENTS

This section describes the functional requirements addressed by the use-cases and specification changes described in all subsequent sections of this document.

5.1 UE Provisioning Requirements

UE Provisioning involves the provisioning of both the eCM and the eUE ('eDVA' and 'eESG' here). This section describes two different provisioning scenarios: The basic IPv6-only provisioning and the Dual Stack provisioning of an Embedded UE (eUE) device in a PacketCable Network.

5.1.1 Review of IPv6-only eUE Provisioning

The eCM provisioning is accomplished using the procedures specified by DOCSIS and eDOCSIS, with additional enhancements. An eCM uses the DOCSIS DHCP process to get the eCM configuration information with certain modifications such as obtaining PacketCable eUE DHCP Server information.

PacketCable defines three provisioning flows for the eUE: Secure, Basic and Hybrid, which essentially follow the same process. Once the eCM has completed provisioning, the eUE is initialized using the eUE DHCP Server Address obtained during eCM provisioning. An eCM contained within an E-DVA or E-ESG passes on information obtained in its DHCP process to the eUE component to determine the IP addressing mode preference for the eUE. Depending on the addressing mode preference, the eUE proceeds to provision in either IPv4 mode or IPv6 mode. The eUE acquires the ToD from the eCM. In the eUE IPv6 Secure Provisioning Flow, the eUE constructs a link-local address for its management interface according to [RFC 4862]. After successful link-local address assignment is accomplished, the eUE performs discovery of the default and neighboring routers by sending Router Solicitation (RS) messages per [RFC 4861]. Once the eUE has completed router discovery, it follows the steps shown in the sequence diagram in Figure 1 which explains the remaining steps in the secure IPv6 provisioning flow. All of these flows use IPv6 and work with back-end servers that support IPv6.

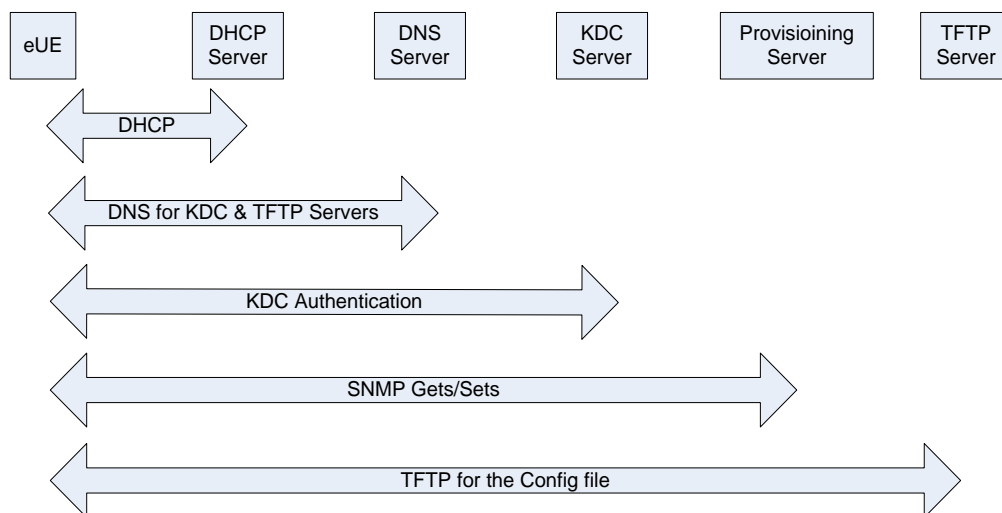


Figure 1 - eUE Provisioning Sequence Diagram

5.1.2 eUE Dual-Stack (IPv4 & IPv6) Provisioning Requirements

eUE provisioning follows the process described in Section 5.1.1. In addition, when configured for Dual Stack mode, it performs DHCP for the second address family, thus acquiring both an IPv4 and IPv6 address. The eUE initializes after the eCM has completed provisioning.

The eUE performs both DHCPv4 and DHCPv6 to acquire addresses from each address family. A Dual-Stack eUE identifies one address family to be "preferred", based on information received from the eCM as described in Section 5.1.1. The eUE will use its "preferred" address family for steps in the provisioning process that should not be replicated, i.e., all signaling and configuration steps after address acquisition (i.e., DHCP). The preferred address flow will follow the normal flows defined currently in the PacketCable specifications while the secondary address flow will include only acquisition of the IP address.

5.1.3 eUE SIP Registration

The eUE client needs to perform SIP registration with the IMS core to enable the voice functionality of the device.

In order for the eUE to register with the IMS Core, it learns the IP address or Fully-Qualified Domain Name (FQDN) of the Proxy Call Session Control Function (P-CSCF) server from its configuration file or SNMP. A Dual-Stack eUE will only register using one address family; the "preferred" address family described in Section 5.1.2.

5.2 Session Signaling and Media Requirements

Once the UE completes the SIP registration process, it is capable of placing a call. In Dual-Stack mode, the media can be carried over either IPv6 or IPv4. The signaling only uses the chosen IP protocol with which it registered.

5.2.1 Dual-Stack Interworking

Although signaling is done over a single version of IP with the IMS core, a Dual-Stack UE must be able to navigate all of the media interworking scenarios listed in Table 1. This is done by indicating what media the UE can support and negotiating the proper IP version for each media flow during call set up. These scenarios must be supported regardless of which version of PacketCable UE-2 is using.

Table 1 - Dual-Stack Media Interworking Scenarios

UE-1	Network-1	Network-2	UE-2
Dual-Stack	Dual-Stack	Dual-Stack	Dual-Stack
Dual-Stack	Dual-Stack	Dual-Stack	IPv6-Only
Dual-Stack	Dual-Stack	Dual-Stack	IPv4-Only
Dual-Stack	Dual-Stack	IPv6-Only	Dual-Stack
Dual-Stack	Dual-Stack	IPv6-Only	IPv6-Only
Dual-Stack	Dual-Stack	IPv4-Only	Dual-Stack
Dual-Stack	Dual-Stack	IPv4-Only	IPv4-Only
Dual-Stack	IPv6-Only	Dual-Stack	Dual-Stack
Dual-Stack	IPv6-Only	Dual-Stack	IPv6-Only
Dual-Stack	IPv6-Only	IPv6-Only	Dual-Stack
Dual-Stack	IPv6-Only	IPv6-Only	IPv6-Only
Dual-Stack	IPv4-Only	Dual-Stack	Dual-Stack
Dual-Stack	IPv4-Only	Dual-Stack	IPv4-Only
Dual-Stack	IPv4-Only	IPv4-Only	Dual-Stack
Dual-Stack	IPv4-Only	IPv4-Only	IPv4-Only
NOTE: Scenarios not listed here may require a Translation Gateway to perform IPv4/IPv6 media interworking in order to function.			

6 PROVISIONING

Prior to this report a PacketCable 2.0 SIP endpoint (e.g., an E-DVA) gains IP connectivity using a single IP stack (IPv4 or IPv6) for its provisioning and functional operations. A dual-stack PacketCable 2.0 SIP endpoint requires both IP stacks to be initialized and configured. This section details the requirements to extend the provisioning interfaces requirements for dual-stack operation.

This section describes the dual stack requirements for the eUE Provisioning Framework (Embedded UE) defined in [ARCH-FRM TR]. The UE Provisioning Framework (Standalone UE) [ARCH-FRM TR] dual stack requirements will be deferred for future study. The main reason for approaching the embedded UE case and not the standalone UE relies on PacketCable 2.0 capabilities to control the boot process of an E-UE to acquire single or dual IP stack, while the standalone box might not provide the appropriate services to perform the same operation.

6.1 Dual Stack Provisioning requirements

The optimal goal of PacketCable dual stack is to provide end-to-end SIP communication establishment between SIP agents without IP media relays or translation gateways (see Section 7). The PacketCable provisioning framework only needs to guarantee the SIP agent obtain IP connectivity in both IPv4 and IPv6 IP stacks which is accomplished in an incremental manner for the current PacketCable 2.0 infrastructure as follows:

- Extend the eUE Provisioning reference point pkt-eue1 [ARCH-FRM TR] between the DHCP server and the eUE to acquire a secondary IP address in the other IP stack as used for Provisioning.
- All existing provisioning processes are executed in a mode referred to as the Primary IP Address Mode. The process of obtaining the secondary IP address is referred to as the Secondary IP Address Mode.

There should be a mechanism to determine which IP address family is used as primary and secondary.

The secondary IP Address mode is subrogated to the success and integrity of the Primary IP address mode provisioning, not the opposite. In other words, if the Primary IP Address mode fails, the Secondary IP Address mode will terminate its current operation. However, a failure in the secondary IP address mode will not affect the operation of the Primary IP Address mode.

All the additional eUE provisioning Reference Points procedures pertaining the eUE (i.e., pkt-eue2, pkt-eue-3, pkt-eue-4, pkt-eue-5, pkt-eue-6) will be executed exclusively through the Primary IP Address mode.

The Primary IP address mode requires a defined completion with either retry mechanisms or shutdown of the eUE requiring manual intervention in case of failure. Contrarily, the Secondary IP Address mode can be configured to continue or stop attempting IP connectivity completion.

6.2 SIP Registration configuration

Discovery of P-SCSF and registration can be performed in two ways:

- Fix IP address Mode, by using either the IPv4 or IPv6 eUE stack to contact the P-CSCF and performs registration;
- Negotiation the selection of the P-CSCF IP registration and signaling plane.

This document indicates requirement for Fix IP address mode SIP operation. In particular, the same IP stack used for provisioning need to be used for SIP signaling. Future versions of this report may introduce mechanisms to control the negotiation or selection of the signaling IP plane when in dual stack. See Section 7 more functional requirements in this regard.

6.3 Preferred Media Advertisement

As part of the configuration server eUE reference point (pkt-eue-5) and based on the interest of traffic, the eUE can be configured to prefer one media IP version over the other when the session establishment procedures indicate that both IP versions are supported by the remote endpoint.

7 SESSION SIGNALING AND MEDIA

Section 6 describes the provisioning processes whereby a PacketCable 2.0 SIP endpoint (e.g., an E-DVA or E-ESG) capable of supporting dual-stack acquires its IPv4 and IPv6 host addresses. This section describes how sessions are established when there is a mix of single and dual-stack SIP entities in the signaling chain between (and including) the originating and terminating endpoints. It describes the mechanisms for interworking between incompatible IP versions at the IP and SIP layer within the SIP signaling plane. It also describes IP version interworking at the IP layer for the media plane, and how the media IP version is negotiated when one or both of the endpoints in a call support dual-stack.

7.1 IP Version Interworking at the IP Layer

7.1.1 IP Layer in SIP Signaling Plane – within PacketCable 2.0 Core

SIP message routing at the IP layer is hop-by-hop in a PacketCable 2.0 network; i.e., endpoints don't send message directly to each other over the IP network, but route the messages via the set of SIP nodes/proxies in the IMS core. Therefore, as long as adjacent SIP nodes in the signaling chain support compatible IP versions, there will be no IP version interworking issues at the IP layer when routing SIP messages.

The routing information for a SIP message can be locally configured in the SIP proxy, or carried in the SIP message itself. The routing information can be in the form of an IPv4 or IPv6 IP address, in which case the SIP proxy uses the address directly to route the message to the next-hop. Or, the routing information can be in the form of an FQDN that the SIP proxy resolves using DNS to an IP address that it then uses to route the message to the next hop.

Therefore, once the IP network infrastructure (e.g., IP routers) supporting a PacketCable 2.0 network has been updated to support dual-stack, the operator can begin to migrate the PacketCable 2.0 core network elements such as the CFCFs and IBCFs to dual-stack. The hop-by-hop routing property of SIP messages avoids IP version interworking issues at the IP layer when the network contains a mix of single-stack IPv4 and dual-stack SIP entities. As long as the IP routing databases accurately reflect the IP version capabilities of the sending and receiving SIP entity, the SIP message will be routed properly. If the sending node is dual-stack, it will always select its host address for SIP signaling that matches the IP version of the next-hop entity.

7.1.2 IP Layer in Signaling Plane – between Endpoints and PacketCable 2.0 Core

As described in Section 6, a dual-stack PacketCable 2.0 endpoint will obtain two host IP addresses (IPv4 and IPv6) during the provisioning process, but it will select only one of them for SIP signaling, based on locally configured data. Essentially, the endpoint is single-stack at its SIP signaling interface. To ensure IP version interworking at the IP layer between the endpoint and the core network, the SIP endpoint must select the P-CSCF address whose IP version matches the IP version the endpoint is configured to use for SIP signaling.

7.1.3 IP Layer in Media Plane

Media RTP and RTCP packets are sent end-to-end between the two media endpoints; i.e., the local media endpoint sets the destination address in the IP header of outgoing RTP packets to the host address of the remote media endpoint, and the message is routed directly via the IP network to the remote endpoint. Due to the end-to-end nature of media packet routing, the IP version of the local and remote media endpoints must match.

If one or both of the local and remote media endpoints are dual-stack, the IP version they use for media is negotiated during session establishment using the procedures described in Section 7.3.2. If the local and remote media endpoints support different IP versions, or if there is an IP network segment between the local and remote endpoints that is incompatible with the selected media IP version, then the PacketCable 2.0 network must insert a media relay device to interwork the IP headers between the two IP versions as described in Section 7.3.1.

7.2 IP Version Interworking at the Signaling Layer

7.2.1 Signaling Layer – SIP Header Fields

IP addresses carried in a header field of a SIP message can potentially cause interworking issues when received by a single-stack SIP entity that doesn't support the address IP version. For example, if a dual-stack S-CSCF uses its IPv6 address as the host-name of a SIP URI in the Via or Record-Route header field of a SIP request, this IPv6 address would be received by all downstream SIP entities in the signaling chain, including single-stack IPv4 SIP entities.

If the required behavior of the downstream IPv4 SIP entity is simply to copy the received IP address into a subsequent request or response (say, when a SIP endpoint copies a received Via or Record-Route header into a response), there should be no interworking issue, since this behavior is required by [RFC 3261]. A network that does encounter interworking issues for this case (say it is serving SIP endpoints that aren't compliant with this aspect of [RFC 3261]) would have to deploy an interworking entity such as an IBCF ALG in the signaling chain to convert the addresses into a form compatible with the non-compliant entity.

NOTE: The above IP version interworking case can occur when the SIP nodes in the network core are updated to support dual-stack even though all the endpoints are still single-stack IPv4.

An interworking issue (problem) would theoretically occur if the downstream IPv4 SIP entity is required to actually use an incompatible IP address to send a subsequent message. Fortunately, this situation should not occur in real deployment scenarios. For example, a SIP endpoint is required to send a SIP response to the top-most address received in the Via header field of the request. By definition, this address should always be supported by the SIP endpoint, since it identifies an adjacent SIP entity in the signaling chain (i.e., the P-CSCF in an IMS architecture), which by design must support a compatible IP version.

7.2.2 Signaling Layer – SDP Body

IP version interworking issues can occur at the SIP layer due to incompatible IP addresses received in the SDP message body. This document describes two mechanisms for resolving this type of interworking issue; the existing PacketCable 2.0 IP version interworking procedures for single-stack endpoints, and the ICE-lite procedures for dual-stack endpoints being newly introduced by this effort. These mechanisms are described in more detail in Section 7.3:

- Section 7.3.1 describes the IP version interworking procedures that are currently supported by PacketCable 2.0 network in a single-stack environment when there is a mix of IPv4-only and IPv6-only endpoints.
- Section 7.3.2 describes how the ICE-lite mechanism defined in [RFC 5245] is used to negotiate the media IP address when one or both endpoints in a session are dual-stack.
- Section 7.3.3 describes how the existing PacketCable 2.0 IP version interworking procedures described in Section 7.3.1 and the new ICE-lite procedures described in Section 7.3.2 work together to provide an overall IP version interworking framework required to migrate a PacketCable 2.0 network from IPv4 to IPv6.

7.3 IP Version Interworking Procedures in PacketCable 2.0

7.3.1 IP Version Interworking in a Single-Stack Environment

This section summarizes IP version interworking mechanisms/procedures described in 3GPP [TS 23.228] and the CableLabs specification [PKT 24.229]. Refer to these specifications for further detail on the mechanisms/procedures.

3GPP [TS 23.228] defines the following mechanisms for IP Version interworking:

- Annex I of 3GPP [TS 23.228] interworking between IMS IPv6 and IPv4 domains using the IMS-ALG function of the IBCF, which supports IP Version interworking capabilities. The IMS-ALG controls the TrGW, which supports media plane aspects of IP Version interworking.
 - IPv6 IMS Domain to IPv4 IMS Domain interworking:

- The originating S-CSCF in the IMS IPv6 domain determines, using DNS or other mechanisms, that the terminating IMS domain supports IPv4 only. The originating S-CSCF forwards the request to the IBCF/IMS-ALG within the originating IMS domain.
- IPv4 IMS Domain to IPv6 IMS Domain interworking:
 - The IMS IPv6 domain provides an IBCF/IMS-ALG as an entry point for IMS IPv4 domains.
- Annex G of 3GPP [TS 23.228] defines the IMS-ALG function within the P-CSCF, as part of the general reference model for IP address translation. The IMS-ALG controls the IMS Access Gateway, which can support media plane aspects of IP Version interworking, between the UE and the backbone network.

[PKT 24.229] includes procedures for IP Version interworking as follows:

- Procedures in support of 3GPP [TS 23.228] Annex G and Annex I mechanisms noted above. Procedures relevant to Annex I are defined within section 5 in [PKT 24.229]. Procedures relevant to Annex G are defined within Annex G in [PKT 24.229]. Note that [PKT 24.229] Annex G defines usage of the P-CSCF controlled IMS Access Gateway for IP version interworking, but does not normatively define triggers/conditions for insertion of the IMS Access Gateway for IP version Interworking.
- In addition, an IBCF/IMS-ALG may also be inserted for IP Version interworking in the following cases, as defined within section 5 of [PKT 24.229]:
 - An originating or terminating S-CSCF that receives an error response to an offer, which indicates that the IP address type in the offer is not supported, may cause an IBCF within the S-CSCF's IMS subsystem to be dynamically inserted.
 - A terminating S-CSCF may examine an SDP offer to determine whether the IP address type is supported within the IMS subsystem and cause an IBCF in the terminating IMS subsystem to be dynamically inserted.
- In addition, as defined within section 6 of [PKT 24.229], an IMS UE that receives an error response to an SDP offer with an IPv6 media address, which indicates that the IPv6 media address is not supported, can send a new INVITE using an IPv4 media address.

7.3.2 Negotiating Media IP Version among Dual-Stack Endpoints

This section describes how the ICE-lite procedures defined in [RFC 5245] are used during session establishment by a dual-stack PacketCable 2.0 endpoint to negotiate an IP version for media that is compatible with the remote endpoint.

The dual-stack PacketCable 2.0 endpoint supporting ICE-Lite should support the following requirements, as defined by [RFC 5245]:

- 1) Only "host" candidates (physical or virtual attached to an interface on the host) are selected as ICE-lite candidates. The endpoint does not gather other candidate addresses.
- 2) Endpoints are not required to initiate connectivity checks. However, per the ICE-lite procedures, they must be able to respond to connectivity checks in order to interwork with full-ICE endpoints.
- 3) RTP-based media has two components; RTP with a component ID of 1, and RTCP with a component ID of 2.
- 4) Endpoints allocate two candidate addresses for each component; one IPv4 and one IPv6. For example, when establishing a voice call, the ICE-lite endpoint must select an IPv4 and IPv6 candidate for the RTP component, and an IPv4 and IPv6 candidate for the RTCP component.
- 5) Endpoints must assign a unique priority for each candidate for the same media stream. Since these candidate priority values only govern the order of the candidate checks when the remote endpoint supports full-ICE, the actual priority values used are outside the scope of this specification effort.
- 6) Once the initial offer/answer exchange has completed, both PC2.0 endpoints examine their own and their peer's candidates. For each media stream, each PC2.0 endpoint pairs up its own candidates with the candidates of its peer for that media stream (same component, UDP, same IP address family).

- 7) When the ICE-lite procedures yield multiple candidate pairs across both IP versions for a media stream, the endpoint must select a valid candidate pair from the preferred media IP version indicated by the configuration parameter described in Section 6.3. In addition, when selecting among multiple IPv6 candidate pairs (say IPv6 is the preferred media IP version), the Endpoint must select a valid candidate pair based on the procedures defined in [RFC 3484].
- 8) In order to enable interworking with single-stack IPv4-only endpoints that don't support ICE-lite, the endpoint should select an IPv4 address for the "default" candidate (the IP address in the SDP "c=" line).
- 9) Endpoints must include an "a=ice-lite" session-level attribute in the SDP.
- 10) The endpoint that generated the offer which started the ICE processing must take the controlling role, and the other endpoint must take the controlled role. These roles persist for a session unless ICE is restarted.
- 11) Endpoints must not send media until they have constructed a valid list of candidates that contains a candidate pair for each component of that media stream.

7.3.2.1 Basic Call Establishment between two Dual-stack Endpoints

Figure 2 shows the session establishment message flow between two dual-stack endpoints when ICE-lite is used to negotiate the IP version for media. (For simplicity, the ACK message acknowledging receipt of the 200OK to INVITE is not shown in the following message-flow diagrams.)

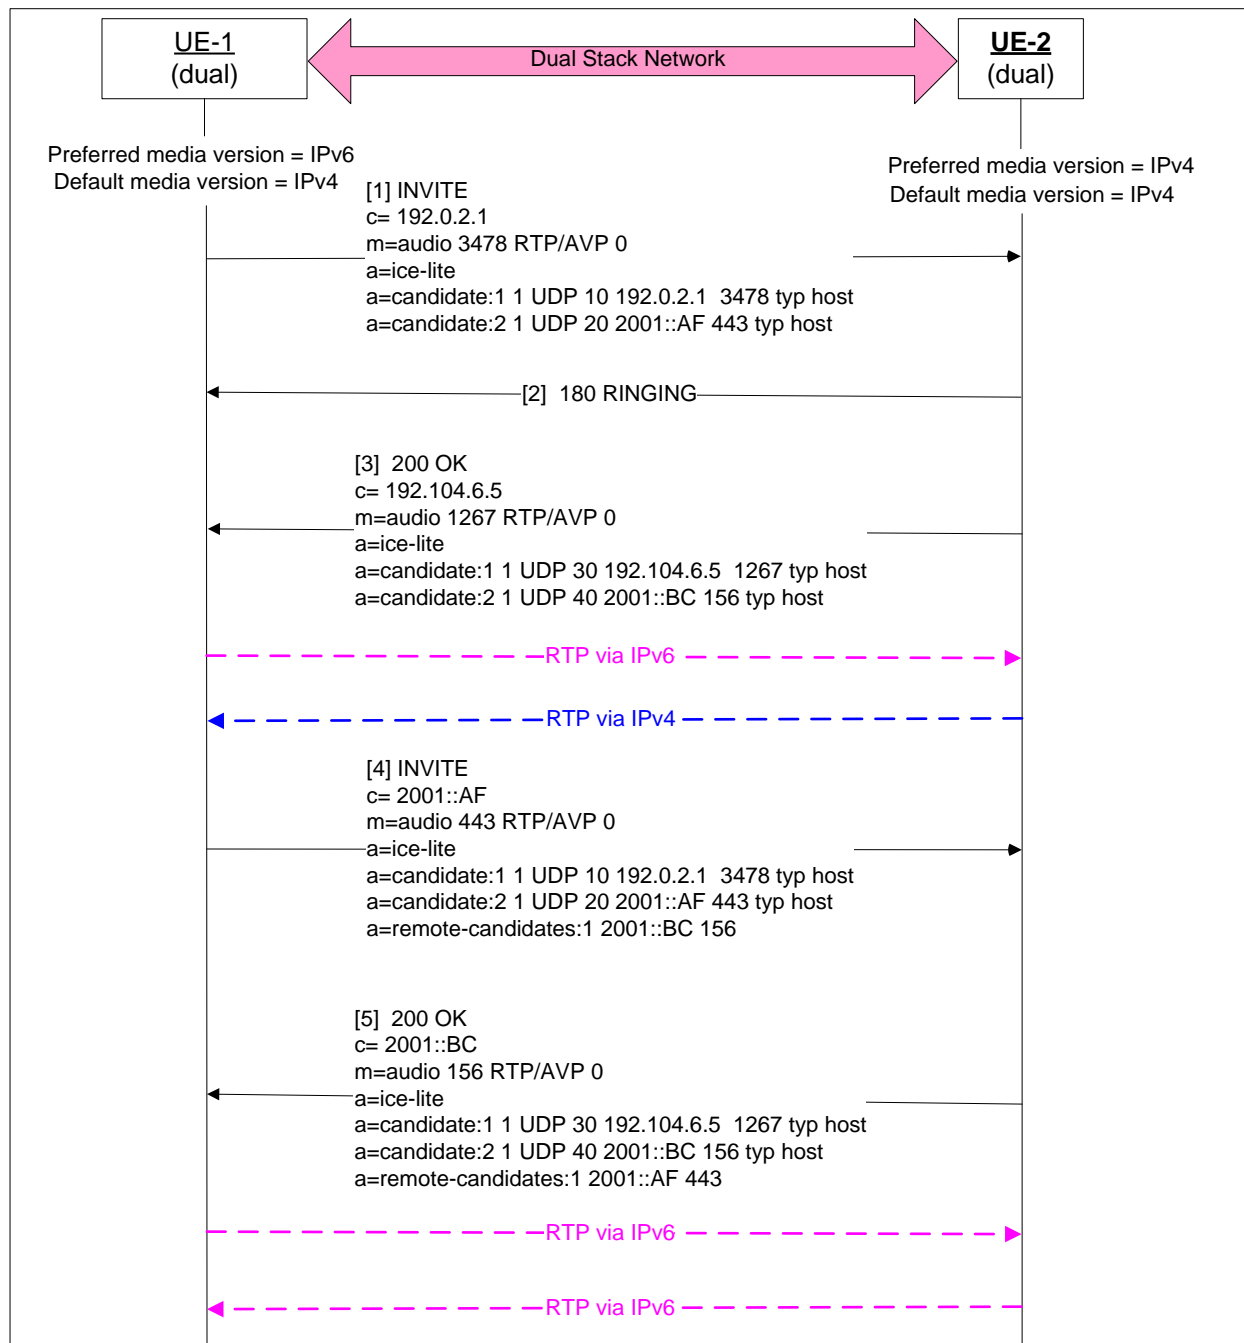


Figure 2 - Using ICE-lite to Negotiate Media IP-version

Initial state:

- Both UE-1 and UE-2 are dual-stack and have obtained a two IP addresses; one IPv4 and one IPv6.
- Both UE-1 and UE-2 are registered via IPv4.
- Both UE-1 and UE-2 are hard coded to use their IPv4 address as the default candidate (as recommended by [RFC 5245]).

- When the ICE-lite candidate negotiation procedure yields multiple candidate pairs:
 - UE-1 is configured to prefer IPv6 for media
 - UE-2 is configured to prefer IPv4 for media

Message Sequence:

[1] UE-1 populates the INVITE SDP-offer as follows:

- The "c=" and "m=" lines contain the default IPv4 address and port respectively, in order to enable interworking with remote single-stack IPv4-only endpoints.
- Two candidate addresses, one for IPv4 and one for IPv6. The candidate parameters shown in the diagram are as follows:

```
"candidate:" <foundation> <component> <transport> <priority> <IP address> <IP
port> "typ" <candidate type, which is always "host" for ICE-lite>
```

How the endpoint selects the candidate priority value is not specified by PacketCable as part of this dual-stack effort, since it only affects the order of candidate checks when the remote endpoint supports full ICE (basically, the priority value is a "don't care" for the PacketCable dual-stack endpoint). Note that normally there are candidates for two components for a single media stream; one for RTP and one for RTCP. For simplicity, this diagram shows only the candidate for the RTP component.

- [2] On receiving [1], UE-2 alerts the called user, and sends a 180 Ringing response per normal SIP procedures. On receiving [2], UE-1 applies local ring-back tone.
- [3] When the call is answered, UE-2 sends a 200OK response to INVITE, containing an SDP-answer populated similar to the SDP-offer in [1] except using UE-2 candidate addresses. Once UE-2 sends [3] and after UE-2 receives [3], both endpoints independently select a valid candidate pair based on their locally configured IP version preference for media identified above; UE-1 selects IPv6 while UE-2 selects IPv4. Per the ICE-lite procedures, both UEs start sending RTP on their selected valid candidate pair, and are prepared to receive RTP on either candidate pair. As a result, the media stream is asymmetrical with-respect-to IP version at this point; the media stream from UE-1 to UE-2 is IPv6, and from UE-2 to UE-1 is IPv4.
- [4] In its role as the ICE-lite controller, UE-1 signals the selected candidate pair to UE-2 by sending an updated SDP-offer. This updated offer differs from the initial SDP-offer as follows:
- The "c=" and "m=" lines identify the selected local candidate, and
 - The SDP contains a "remote-candidates" attribute identifying the remote candidate of the selected candidate pair.
- [5] On receiving [5], UE-2 stops sending RTP on IPv4 and starts sending on IPv6. On receiving [5], UE-1 stops listening for RTP on IPv4.

7.3.2.2 Basic Call Establishment with Early-Media

Figure 3 shows the session establishment message flow between two dual-stack endpoints when ICE-lite is used to negotiate the IP version for media, and the terminating UE has a feature that requires early-media.

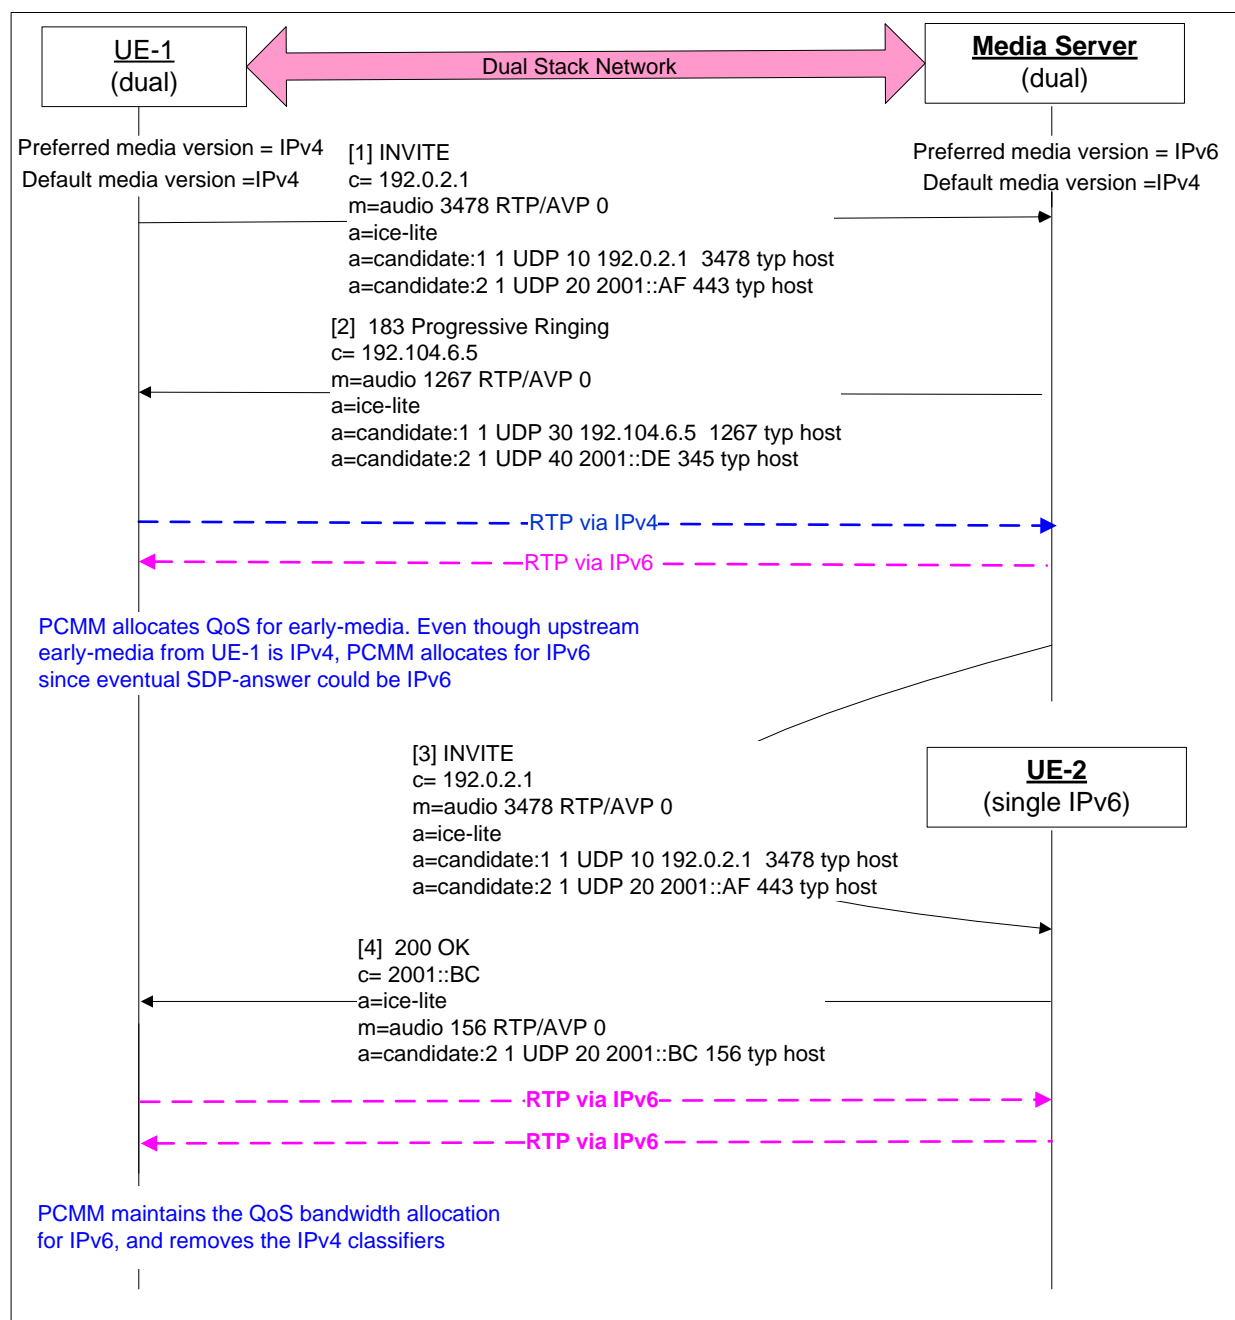


Figure 3 – Support of early media with ICE-lite

Initial state:

- UE-1
 - Dual-stack with two IP addresses; one IPv4 and one IPv6.
 - Default candidate hard-coded to IPv4
 - Preferred media IP version configured to IPv6

- UE-2
 - Single stack IPv6 (since it only has one candidate, by definition the default and preferred IP media candidate is IPv6).
 - UE-2 has a terminating feature (e.g., solicitor call blocking) that requires early-media from a dual-stack Media Server that is configured to prefer its IPv6 media candidate address.

Message Sequence:

UE-1 sends an initial offer in [1]. Since UE-2 has a terminating feature that requires early media, the terminating network routes the INVITE to a Media Server. The Media Server establishes the early-media session by populating an SDP-preview in the [2] 183 response, containing the Media Server's IPv4 and IPv6 candidate addresses.

On receiving [3], UE-1 selects a valid candidate pair based on its configured preference; in this case IPv4. Likewise, the Media Server selects the IPv6 candidate pair. Both endpoints start sending on their selected valid candidate pair, and listening on both candidate pairs, as dictated by ICE-lite; i.e., the early-media session is asymmetrical with-respect-to IP version in the send/receive direction.

Once the Media Server has performed its role for the terminating feature (e.g., prompted for and collected a solicitor-call-blocking PIN code from the calling user), the Media Server sequentially forks the initial [1] INVITE to UE-2 (i.e., INVITE [1] and [3] have same SDP-offer, and contain no To-tag).

On receiving [3], UE-2 alerts the called user. At this point in the call the calling user should hear ring-back tone. This use-case doesn't show the details of how that would occur, but one way would be for the Media Server to send ring-back tone RTP via the early-media session.

When the called user answers the call, UE-2 sends [4] 200OK response containing an SDP-answer. Since UE-2 is single-stack, the SDP-answer contains only a single IPv6 candidate.

NOTE: The [4] 200OK response will have a different To tag than the [2] 183 response, and therefore will establish a separate dialog. Also, to minimize complexity, the 180 Ringing response from UE-2 is not shown.

Since there is only a single candidate pair, candidate selection is trivial; UE-1 and UE-2 both select and use the single IPv6 candidate pair. Also, UE-1 is not required to send an updated SDP-offer to signal which candidate-pair was selected.

7.3.3 IPv4/6 Interworking in a Mixed Single/Dual-Stack Environment

To support session establishment in a network that has a mix of single stack IPv4 or IPv6, and dual-stack endpoints, both of the mechanisms described previously are required. The network inserts media relay devices (e.g., IBCF/ALG-TrGW) to enable interworking between incompatible single-stack endpoints as described in Section 7.3.1. If one or both endpoints in a session support dual-stack, then they can negotiate a common media IP version (and therefore eliminate the need for a media-relay device) using the ICE-lite as described in Section 7.3.2. This is the primary advantage of adding support for dual-stack – it minimizes the number of use-cases where insertion of an expensive media-relay device is required.

As described in Section 7.3.1, an IBCF/ALG-TrGW can be inserted statically for all calls, or dynamically per call.

7.3.3.1 ICE-lite with Static IBCF/ALG-TrGW

As described in Section 7.3.1, the IBCF/ALG-TrGW media-relay device may, based on local policy, be inserted statically for all calls; say to interwork between a single-stack IPv6 network and a single-stack IPv4 network (in this case the two networks could be two different cable operator networks, or two different segments within a single cable operator network). In addition, local policy may dictate insertion of a IBCF/ALG-TrGW for other reasons; say to perform network topology hiding at a peering interface. In either case, how does the introduction of dual-stack and ICE-lite impact the deployment of the IBCF/ALG-TrGW?

If the IBCF/ALG-TrGW is being inserted statically for IP version interworking purposes, and one of the networks (say the single-stack IPv4 network) is updated to support dual-stack and ICE-lite, then presumably the insertion of the IBCF/ALG-TrGW can be changed from static to dynamic; i.e., take advantage of the ability of ICE-lite to

negotiate compatible end-to-end IP version for media, and insert the IBCF/ALG-TrGW dynamically only for the subset of calls between two incompatible single-stack endpoints.

If local policy dictates static insertion of the IBCF/ALG-TrGW for reasons other than IP version interworking, then it is very likely that the policy decision isn't affected by the introduction of dual-stack and ICE-lite (i.e., the ALG-TrGW would continue to be statically inserted for all calls after dual-stack and ICE-lite were introduced). The question then arises, how does the ALG-TrGW function interact with the ICE-lite function? Say for example, two operators are in the process of migrating their endpoints to support dual-stack/ICE-lite, and they've deployed IBCF/ALG-TrGW at their egress/ingress interconnect interfaces for topology hiding purposes. In this case it would be advantageous if the ALG-TrGW also supported dual-stack and ICE-lite, for the following reasons:

- When connected to a single-stack endpoint, it could adjust the IP version of its media endpoint on that call leg to match the IP version supported by the endpoint, thus eliminating the need to insert a 2nd media-relay device between the ALG-TrGW and the media endpoint.
- When connected to a dual-stack device, it could enable negotiation of a single IP version for the end-to-end media session, which in turn would minimize the media interworking load on the TrGW, and therefore improve TrGW scalability (see call flow in Section 7.3.3.3).

7.3.3.2 ICE-lite with Dynamic IBCF/ALG-TrGW

The IP version interworking procedures in Section 7.3.1 support the case where the IBCF/ALG-TrGW is inserted dynamically when it's needed to perform IP version interworking. For example, if two peering networks support a mix of single-stack IPv4 and IPv6 endpoints, then the ALG-TrGW can be dynamically inserted for IP version interworking only for calls between an IPv4-only and IPv6-only endpoint (i.e., when a call fails with a 488 (Not Acceptable Here) response). If the operators initiate the migration process to dual-stack/ICE-lite, it has no impact on the operation of the IBCF/ALG-TrGW. It will continue to be inserted to perform IP version interworking only when a session is established between two single-stack endpoints with incompatible IP versions.

7.3.3.3 Basic Call Establishment via an ALG/Media-Relay Device

Figure 4 shows the session establishment message flow to establish a media session between two endpoints via an ALG/Media-Relay device such as the IBCF/ALG-TrGW.

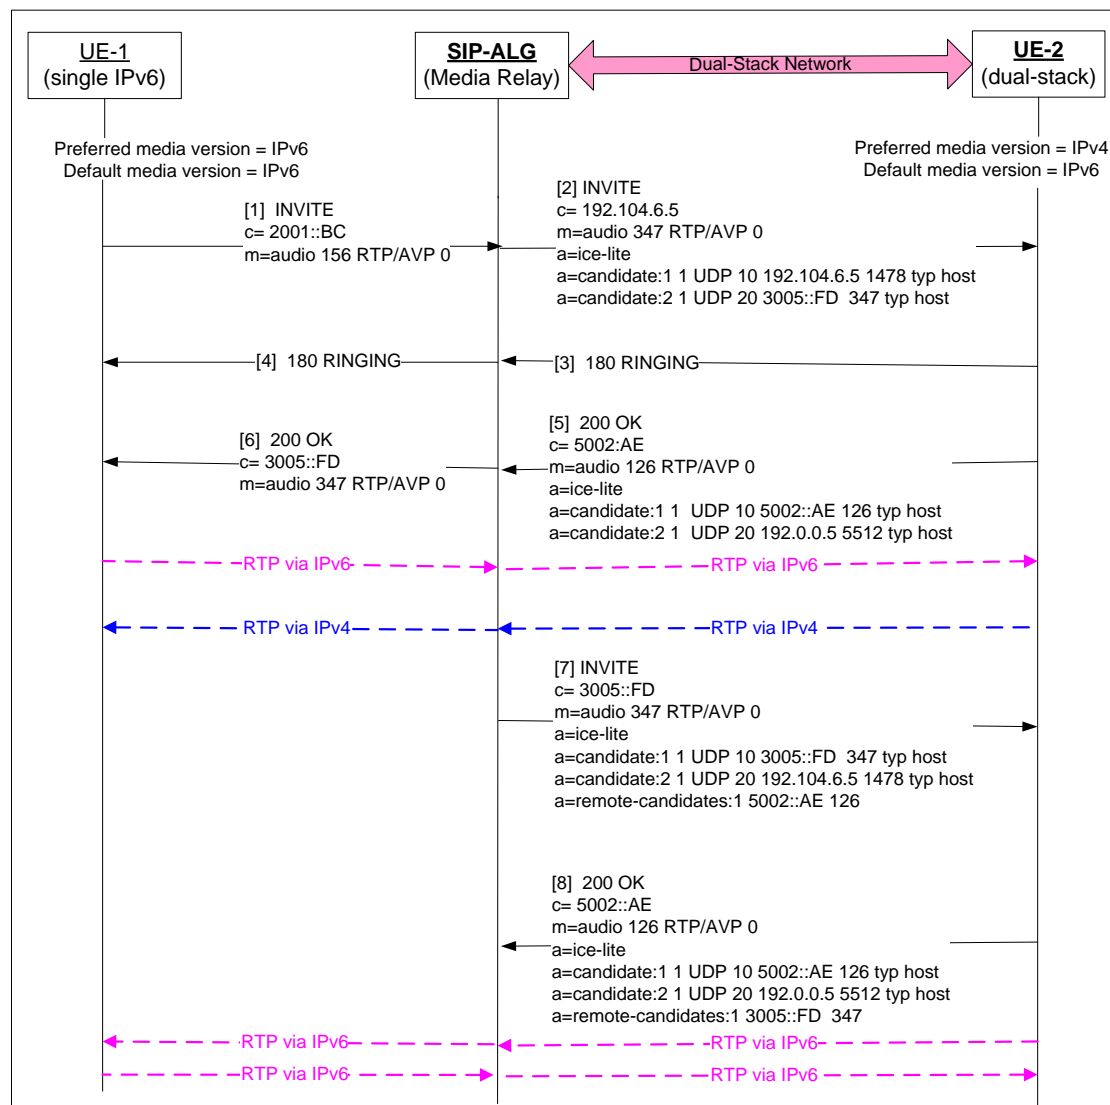


Figure 4 – ALG/Media-Relay Support of ICE-lite

Initial state:

- UE-1
 - Single-stack IPv6.
- UE-2
 - Dual-Stack
 - Default candidate hard-coded to IPv6 (a default of IPv4 is recommended, but IPv6 is being shown here to illustrate how ICE-lite handles this case)
 - Preferred media IP version configured to IPv6

- ALG
- Dual-Stack supporting ICE-lite
 - ALG s statically inserted for all calls

Message Sequence:

Since UE-1 is a single-stack device, the session establishment procedure for the originating call-leg between UE-1 and the ALG follows the normal SDP offer/answer procedures (i.e., does not use ICE-lite). However, the ALG does use ICE-lite for the terminating call leg, so that it can attempt to negotiate the same IP version for the terminating leg that is being used on the originating leg. Having both call legs on the same IP version reduces the media interworking load on the ALG.

On receiving [1], the ALG learns that UE-1 is a single-stack device that supports only IPv6. Therefore, it performs the ICE-lite procedures in the terminating leg with a preference for the IPv6 candidate pair. The ALG populates the SDP-offer in [2] with its IPv4 and IPv6 candidate addresses. When it receives the SDP-answer in [5], the ALG selects a valid candidate pair with a preference for IPv6. (Note that the "c=" line in [5] contains the default IPv6 candidate address.) Since UE-2 is also dual-stack, the ALG is able to select the IPv6 candidate pair. UE-2 has a preference for IPv4, and so selects the IPv4 candidate pair. At this point the media IP version is asymmetrical for the terminating call leg, and therefore the ALG must perform IP version interworking for the RTP stream from UE-1 to UE-1.

In its role as ICE-lite controller for the terminating call leg, the ALG signals the selected candidate pair for the terminating leg to UE-2 by sending an updated SDP-offer in [7]. On receiving [7], UE-2 stops sending RTP on IPv4 and starts sending on IPv6, resulting in IPv6 being used end-to-end for the media stream in both directions.

7.4 Quality of Service and PCMM

PCMM sees all the SDP offer/answer exchanges between the originating and terminating endpoints, and allocates QoS to support the session being established. On an initial offer/answer to establish a session between two dual-stack endpoints, PCMM doesn't know whether the IPv4 or IPv6 candidate pair will be selected. Therefore, it must allocate the QoS bandwidth to accommodate the largest RTP packet size (i.e., the packet with the IPv6 IP header) and assign classifiers to that QoS resource for both candidates. Once the controlling endpoint selects a valid candidate pair, it sends an updated offer/answer to signal the selected candidate pair to the remote endpoint as described in section 7.3.2. As a result of this exchange, PCMM learns the selected candidate pair. It removes the classifiers associated with the candidate that was not selected, and adjusts the QoS allocation downward if the IPv4 candidate pair was selected.

NOTE: The additional offer/answer mandated by ICE-lite to signal the selected candidate will increase the per-call transaction load on the PCMM infrastructure as the number of dual-stack endpoints served by the network increases. There may be ways to reduce this additional transaction load. For example, if an operator knows that all of its deployed dual-stack endpoints prefer IPv4, then the operator could apply a rule at the Policy Server so that if the initial offer/answer contains both an IPv4 and IPv6 candidate pairs, then QoS is allocated for the IPv4 candidate pair. We could then relax the requirement to send the updated offer/answer exchange. (This optimization isn't supported by the PacketCable dual-stack requirements, which mandate support for ICE-lite which in turn mandates the support for the updated offer/answer.)

8 USE CASES

This section describes the SIP message flows to support a number of different IP version interworking use case that would typically be encountered as operators migrate their networks from IPv4 to IPv6.

8.1 Ideal IPv4-to-IPv6 Migration Strategy

This section describes an "ideal" and orderly migration path from IPv4 to IPv6 for a PacketCable 2.0 network. Operators may not be able to follow this ideal path exactly due to real-world considerations, such as a subset of the network endpoints cannot be updated to dual-stack. However, describing the ideal flow will serve to identify the use-cases that the overall solution needs to support.

The ideal migration path consists of four steps:

Step 1 - Migrate the IP core network to dual-stack (e.g., Layer-3 IP routers).

Step 2 - Migrate the PacketCable 2.0 core network to dual-stack (e.g., CSCFs, IBCF).

Step 3 - Migrate the PacketCable 2.0 endpoints to dual-stack (e.g., E-DVA, S-ESG).

Step 4 - Once all endpoints have been upgraded to dual-stack, start deploying IPv6-only endpoints.

At this point PacketCable 2.0 network is fully converted to IPv6 in the sense that all SIP signaling and media packets are being routed using IPv6. The operator can now start to deprecate the IPv4 network infrastructure.

The following sections describe these steps in more detail.

8.1.1 Migrate the IP Core Network to Dual-Stack

At this initial stage, the IP core network is upgraded to support dual-stack, but all the PacketCable 2.0 SIP nodes, including core network elements and endpoints, are single-stack IPv4. All SIP signaling and media packets are routed using IPv4, with the result that there are no IPv4/6 interworking use cases.




Figure 5 - Dual-Stack IP Network

8.1.2 Migrate the PacketCable 2.0 Core Network to Dual-Stack

The SIP nodes in the core network are upgraded to dual-stack. The network could be updated in segments, in which case there would be a mix of IPv4-only and dual-stack network segments. Even though the endpoints are still single-stack IPv4, the operator could update the SIP routing database (e.g., DNS) such that adjacent core SIP-nodes could signal over IPv6, and therefore could add IPv6 addresses to the SIP header fields. The IPv4 endpoints would receive and be expected to handle these IPv6 addresses. At this stage all media traffic is routed via IPv4.




Figure 6 - Dual-Stack PacketCable 2.0 Network

IP version interworking at the IP layer happens as part of normal SIP routing; i.e., a dual-stack SIP proxy can receive a SIP message at one IP version, and then resolve the next-hop address to send the message on a different IP version as described in Section 7.1.1.

IP version interworking at the SIP layer may be required at this stage, based on the capabilities of the IPv4-only endpoints. We expect that most [RFC 3261]-compliant single-stack IPv4 endpoints will correctly handle any IPv6 addresses received in SIP header fields (see Section 7.2.1 in which case no special interworking procedures are required. If the network serves single-stack IPv4 endpoints that cannot handle IPv6 addresses received in SIP header fields, then the operator would have to deploy IBCF ALGs to perform IPv4/6 interworking at the SIP layer.

8.1.3 Migrate the PacketCable 2.0 Endpoints to Dual-Stack

The endpoints are upgraded, over time, to support both dual-stack and the associated ICE-lite procedures to negotiate the IP version used for media. Eventually all PacketCable 2.0 endpoints are converted to dual-stack. Some of the dual-stack endpoints could register using IPv6, in which case the single-stack IPv4 endpoints would start to see additional IPv6 addresses in the SIP messages (e.g., in the Contact header field).

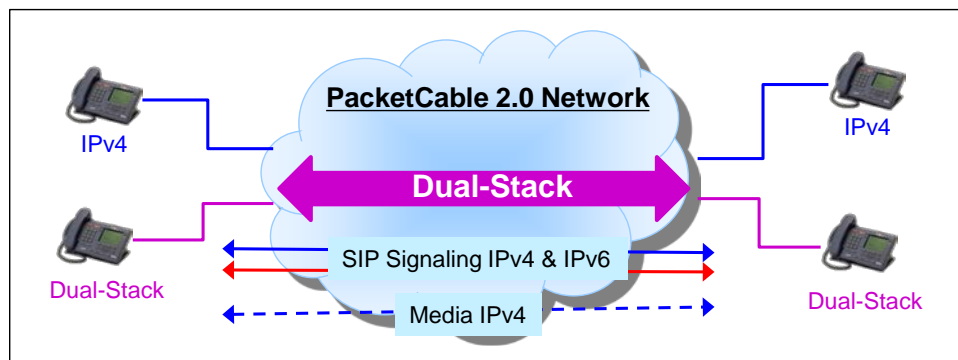


Figure 7 - Dual-Stack Network with Mix of Dual-Stack and IPv4-only Endpoints

The dual-stack endpoints will advertise both IPv4 and IPv6 media addresses in the SDP offer/answer, but since they are configured to prefer IPv4, they will always choose the IPv4 address. Therefore, at this stage, all media packets are all routed via IPv4.

8.1.4 Deploy Single-Stack IPv6-only Endpoints

Once all UEs have been upgraded to dual-stack, the operator can start deploying IPv6-only endpoints. On calls where one or both endpoints are single-stack IPv6 will route media using IPv6. The operator can reconfigure the dual-stack endpoints to prefer IPv6 for media (instead of IPv4 in step-2), in which case media packets will be routed using IPv6.




Figure 8 - Dual-Stack Network with Mix of Dual-Stack and IPv6-only Endpoints

Once all endpoints have been converted to single-stack IPv6, the operator can deprecate the IPv4 infrastructure.

8.2 Deviations from the Ideal IPv4-to-IPv6 Migration Path

Real-world deployments will likely need to deviate from the ideal IPv4→6 migration process described in Section 8.1 for a variety of reasons including:

- Some PacketCable 2.0 endpoints may not be upgradeable to dual-stack.
- Operators will follow independent migration schedules, so that an operator that has completed the migration must be capable of interworking with peering partners that have not started or are still migrating to IPv6.
- PacketCable 2.0 networks must interwork with single-stack IPv4 PacketCable 1.5 networks.

Given these real-world factors, the actual use-case scenarios that will likely be encountered in real-world deployments include the following:

- A PacketCable2.0 network is dual-stack (end-to-end), serving a mix of IPv4-only, IPv6-only, and dual-stack endpoints. Call-flows for this case are described in Section 8.3.1.
- A PacketCable 2.0 network consists of a dual-stack segment and a signal-stack segment (e.g., originating dual-stack network with a mix of single and dual-stack endpoints initiates call to a single-stack IPv4-only terminating network). Call-flows for this case are described in Section 8.3.2.
- A PacketCable 2.0 dual-stack network initiates/receives calls to/from a PacketCable 1.5 network. This case is described in Section 8.3.3.
- Two dual-stack PacketCable 2.0 networks are interconnected by a single-stack IPv4 network. This case is described in Section 8.3.4.

8.3 Call Flow Diagrams

Figure 9 shows the IMS half-call model with the separate originating and terminating networks, and the SIP message components described in the message flow diagrams in this section.

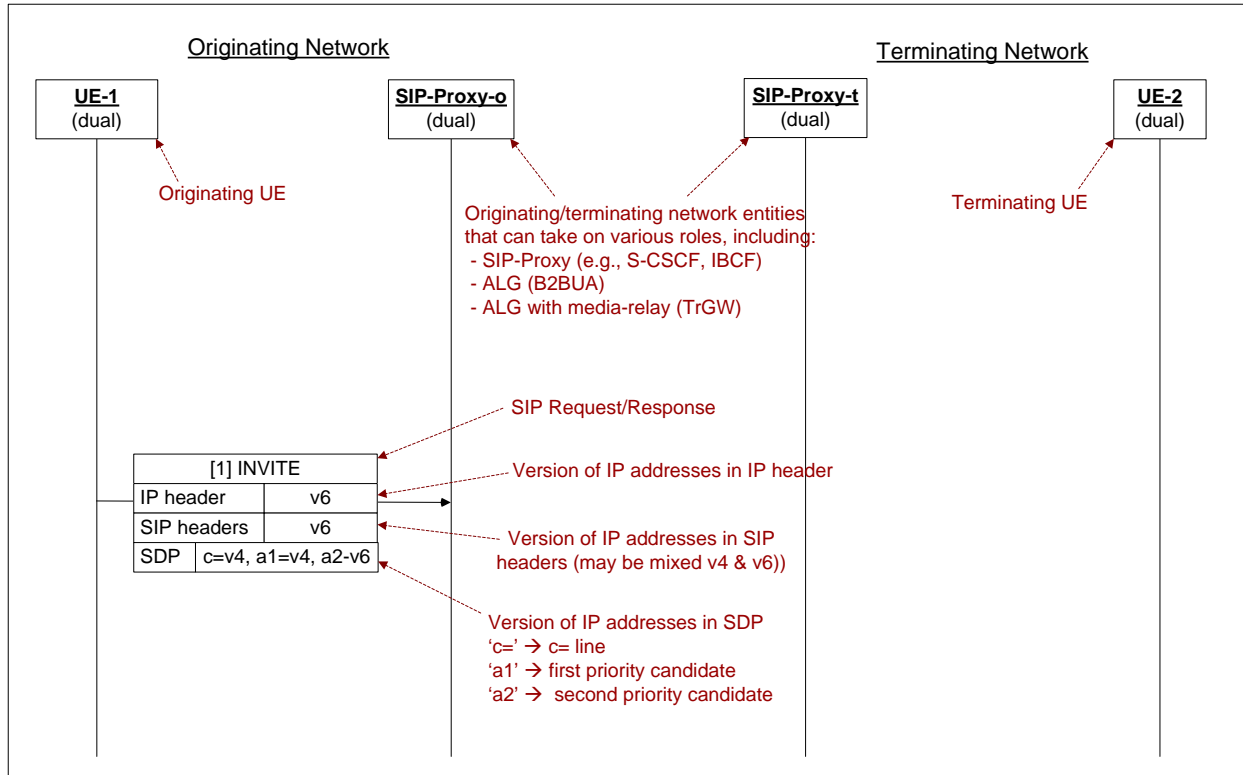


Figure 9 - Message Flow Template

8.3.1 Establishing Calls over a Dual-Stack PacketCable 2.0 Network

This section describes the mainline case where an operator has upgraded the core network to dual-stack, and is in the process of migrating the endpoints to dual-stack, such that the network has a mix of single-stack (IPv4-only or IPv6-only) and dual-stack endpoints. The dual-stack endpoints are registered either via IPv4 or IPv6.

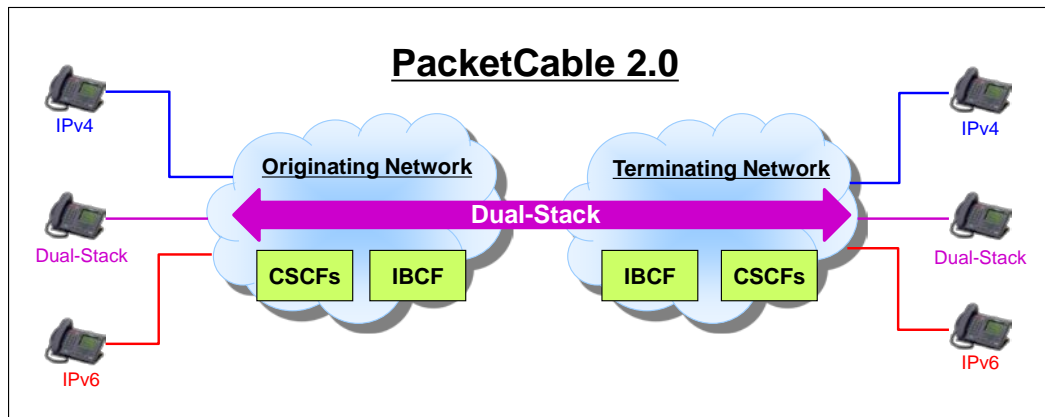


Figure 10 - End-to-End Dual-Stack Network Supporting Mix of Single and Dual-Stack UEs

The following subsections show how calls are established between the different combinations of single-stack IPv4-only and dual-stack endpoints.

8.3.1.1 IP Version Interworking within IP Header and SIP Header over Dual-Stack Network

Figure 11 shows how the IP version is updated within the IP header and SIP header fields as a SIP message traverses the dual-stack network between two single-stack endpoints.

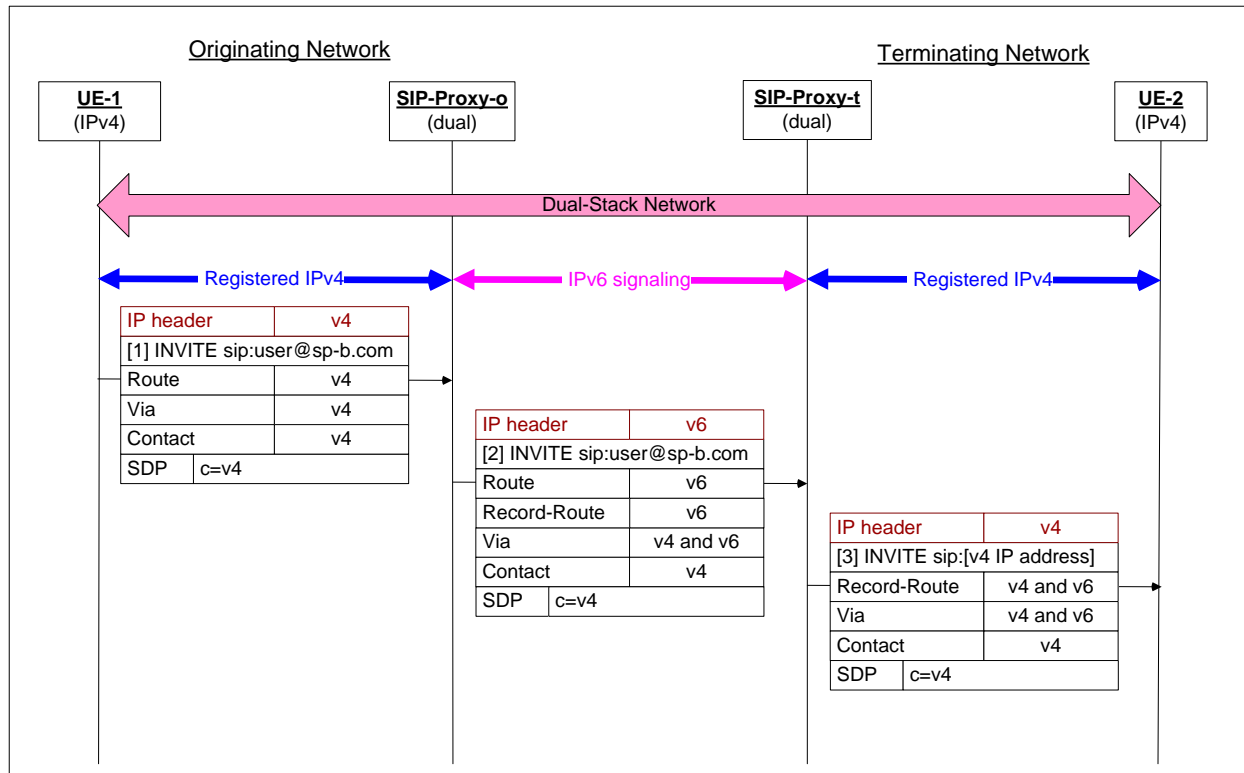


Figure 11 - IP Version Interworking Between Single-Stack UEs in a Dual-Stack Network

Initial state:

- UE-1 and UE-2 are both registered via IPv4.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv6.

Message Sequence:

- [1] UE-1 populates the INVITE request with IPv4 address in the IP header, and in the SIP header fields. Some SIP header fields that could possibly contain IP addresses include:
 - The Route header field containing the service-route discovered during registration.
 - The Via header field containing the address the endpoint wishes to receive subsequent responses to this request.
 - The Contact header field containing the address the endpoint wishes to receive subsequent in-dialog requests.
- [2] Through normal SIP routing logic, SIP-Proxy-o determines that it must route incoming request [1] to SIP-Proxy-t (e.g., resolve the E.164 number identified in the user-part of the Request-URI to the address of the next-hop proxy using DNS or other routing databases). Since the IP address of the next-hop proxy is IPv6, SIP-Proxy-o uses its IPv6 host address as the source address at the IP layer of [2]. Before sending the request, SIP-Proxy-o could add IPv6 addresses to some of the SIP header fields. For example, it could add its IPv6 address to the Record-Route and Via header fields to identify where it wishes to receive subsequent in-dialog requests and responses.

- [3] Again through normal SIP routing logic (e.g., registrar/location database query), SIP-Proxy-t determines that it must forward the request to UE-2. Since the IP address of the next-hop entity is IPv4, SIP-Proxy-t uses its IPv4 host address as the source address at the IP layer of message [3]. Per normal SIP Proxy behavior, SIP-Proxy-t would add its IPv4 address to the received Record-Route and Via header fields.

On receiving [3], UE-2 must copy the received Record-Route and Via header fields, along with any IPv6 addresses they might contain, into the subsequent response. It must also save the route-set contained in the Record-Route header field, so it can use it to build the Route header field for subsequent in-dialog requests.

8.3.1.2 Calls Between Single-Stack IPv4 UEs

Figure 12 shows the session establishment signaling-flow between two single-stack IPv4 endpoints over a dual-stack network (network is dual-stack at IP and SIP layer).

NOTE: For simplicity, the message flow diagrams in Section 8 don't show all aspects of the ICE-lite signaling procedures. For example, the diagrams don't show the subsequent SDP offer/answer exchange that occurs after the call is answered to signal the selected candidate pair between two dual-stack endpoints. For a more detailed description of the ICE-lite procedures, please refer to Section 7.3.2.

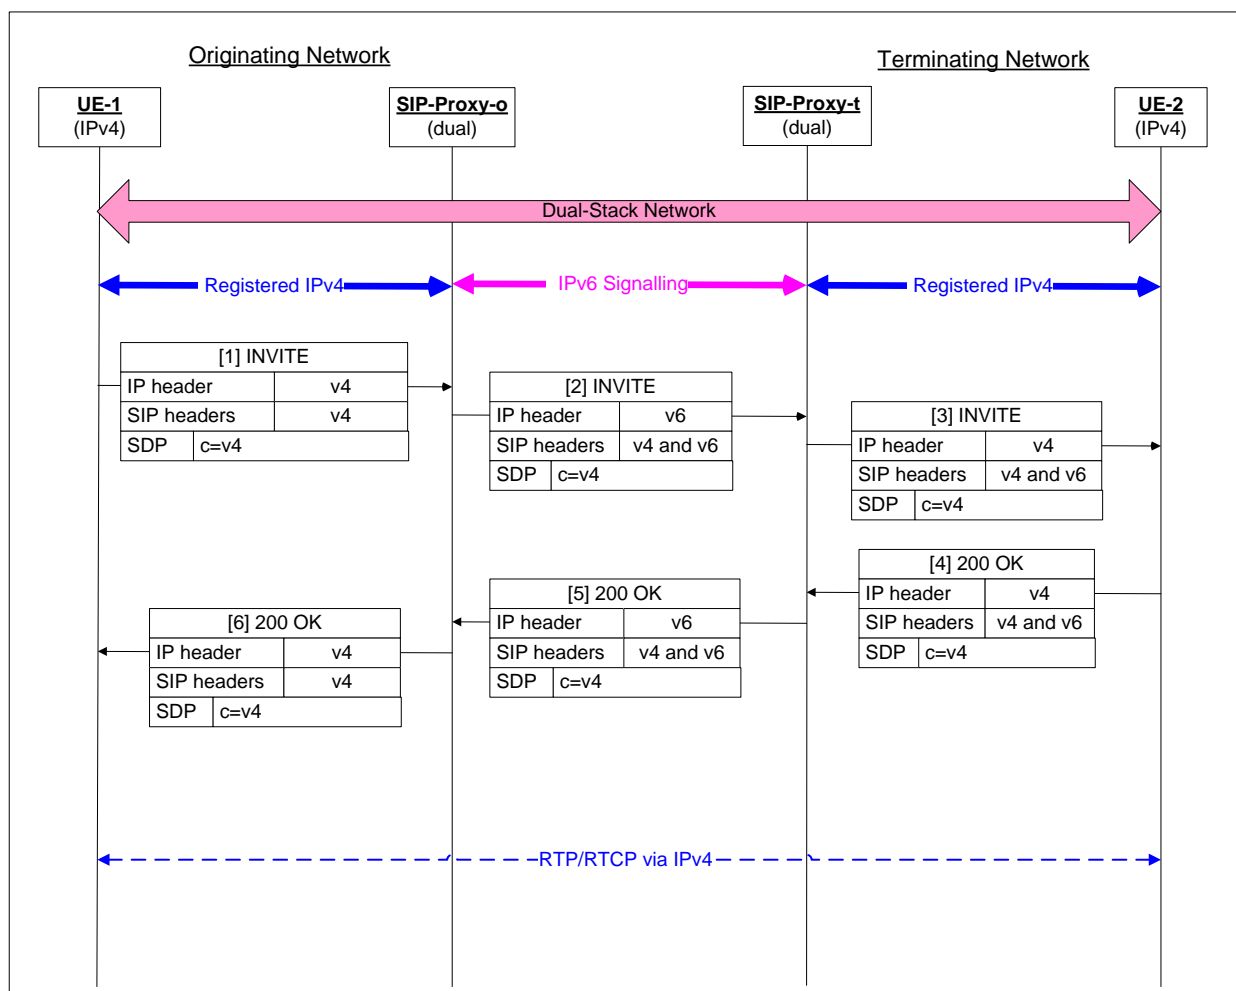


Figure 12 - Single-Stack UE-1 Calls Single-Stack UE-2

Initial state:

- UE-1 and UE-2 are both single-stack IPv4.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv6.

Message Sequence:

- [1] UE-1 populates the INVITE request with IPv4 address in the IP header, in the SIP header fields, and places its IPv4 address in the "c=" line of the SDP offer.
- [2] Before sending the request, SIP-Proxy-o may add IPv6 addresses to some of the SIP header fields.
- [3] SIP-Proxy-t forwards request to UE-2 via IPv4 at the IP layer. Per normal SIP Proxy behavior, SIP-Proxy-t would add its IPv4 address to the received Record-Route and Via header fields, but it will not modify any of the received IP addresses.
- [4] UE-2 returns the SDP-answer containing its IPv4 address in the "c=" line in the 200OK response to INVITE (UE-2 may copy header fields containing IPv6 addresses from [3] to [4], but will otherwise ignore the IPv6 addresses). The SDP-answer is forwarded to UE-1 via [5] and [6], and the media session is established via IPv4.

8.3.1.3 Dual-Stack UEs Registered with Different IP Versions

Figure 13 shows the session establishment signaling-flow between two dual-stack endpoints over a dual-stack network. The endpoints are registered on different IP versions.

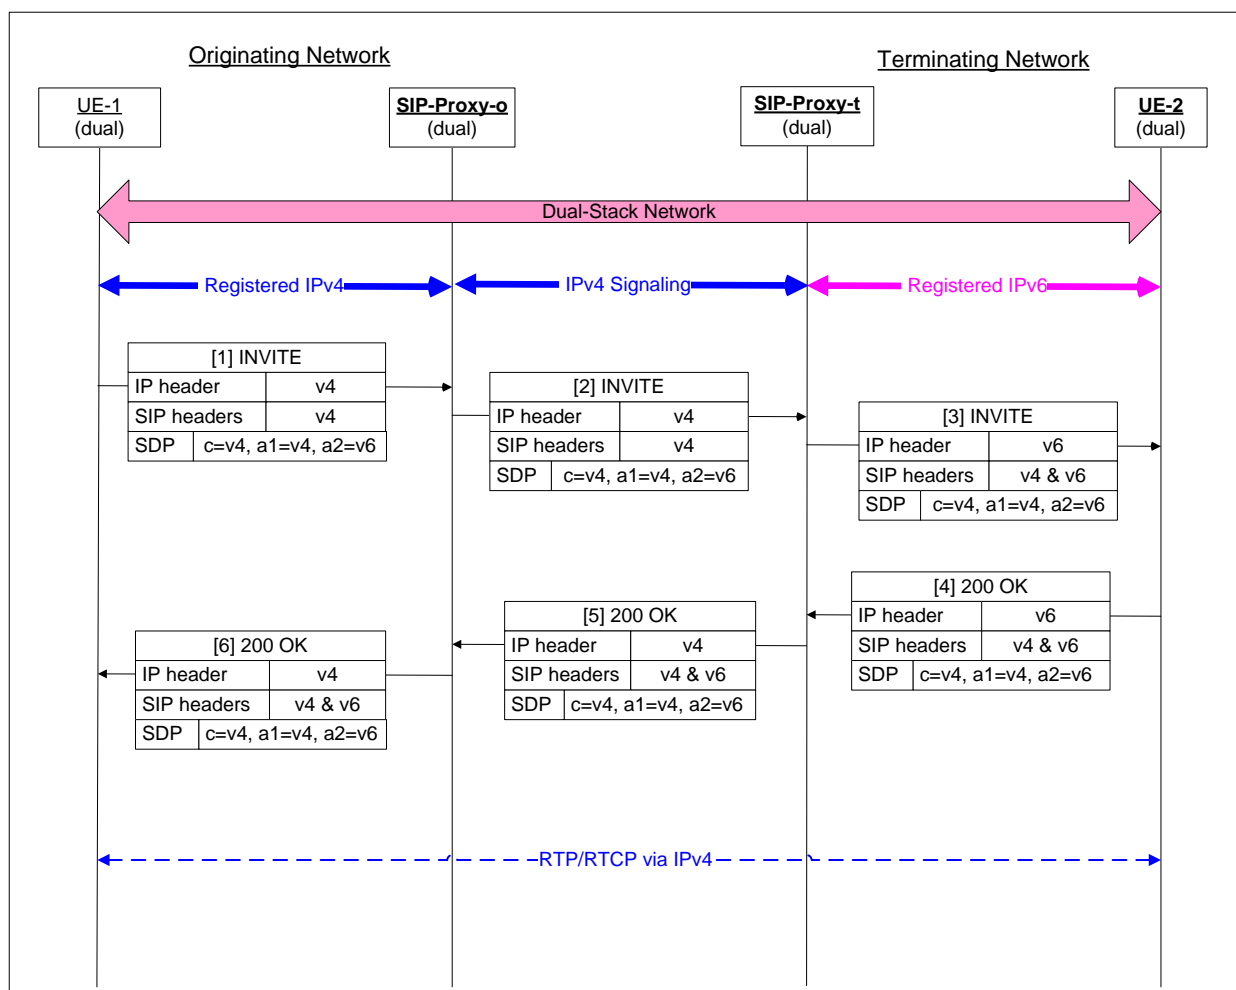


Figure 13 - Dual-Stack UE-1 Registered using IPv4 Calls Dual-Stack UE-2 Registered using IPv6

Initial state:

- UE-1 and UE-2 are both dual-stack, registered via IPv4 and IPv6 respectively.
- UE-1 and UE-2 are both configured to prefer IPv4 for media.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv4.

Message Sequence:

Since both endpoints are dual-stack, the SDP offer/answer contains the ICE-lite candidate addresses advertising both media addresses. Since the endpoints are configured to prefer IPv4 for media, the session is established over IPv4.

8.3.1.4 Calls Between Single-Stack and Dual-Stack UEs

Figure 14 shows the session establishment signaling-flow between a dual-stack and single-stack IPv4 endpoint over a dual-stack network.

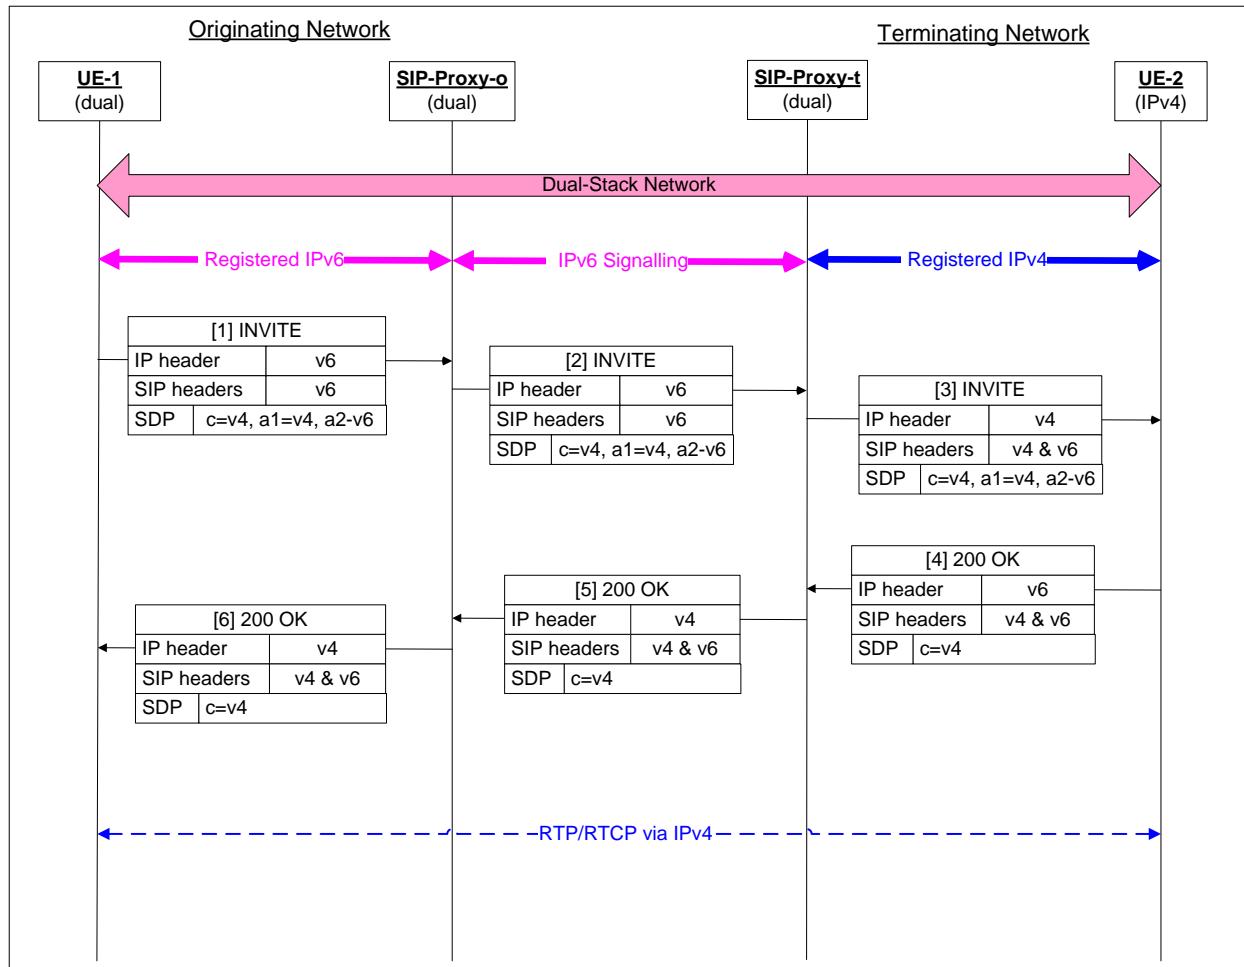


Figure 14 - Dual-Stack UE Registered Using IPv6 Calls Single-Stack IPv4 UE

Initial state:

- UE-1 is dual-stack registered via IPv6 while UE-2 is single stack IPv4.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv4.

Message Sequence:

UE-1 includes uses ICE-lite attributes to advertise both of its IP addresses in the SDP offer. Since UE-2 doesn't support ICE-lite, it ignores the ICE-lite attributes, and sends an SDP-answer corresponding to the IPv4 address received in the "c=" line of the offer, and the session is established over IPv4.

Figure 15 shows the session establishment signaling-flow between a dual-stack and single-stack IPv6 endpoint over a dual-stack network.

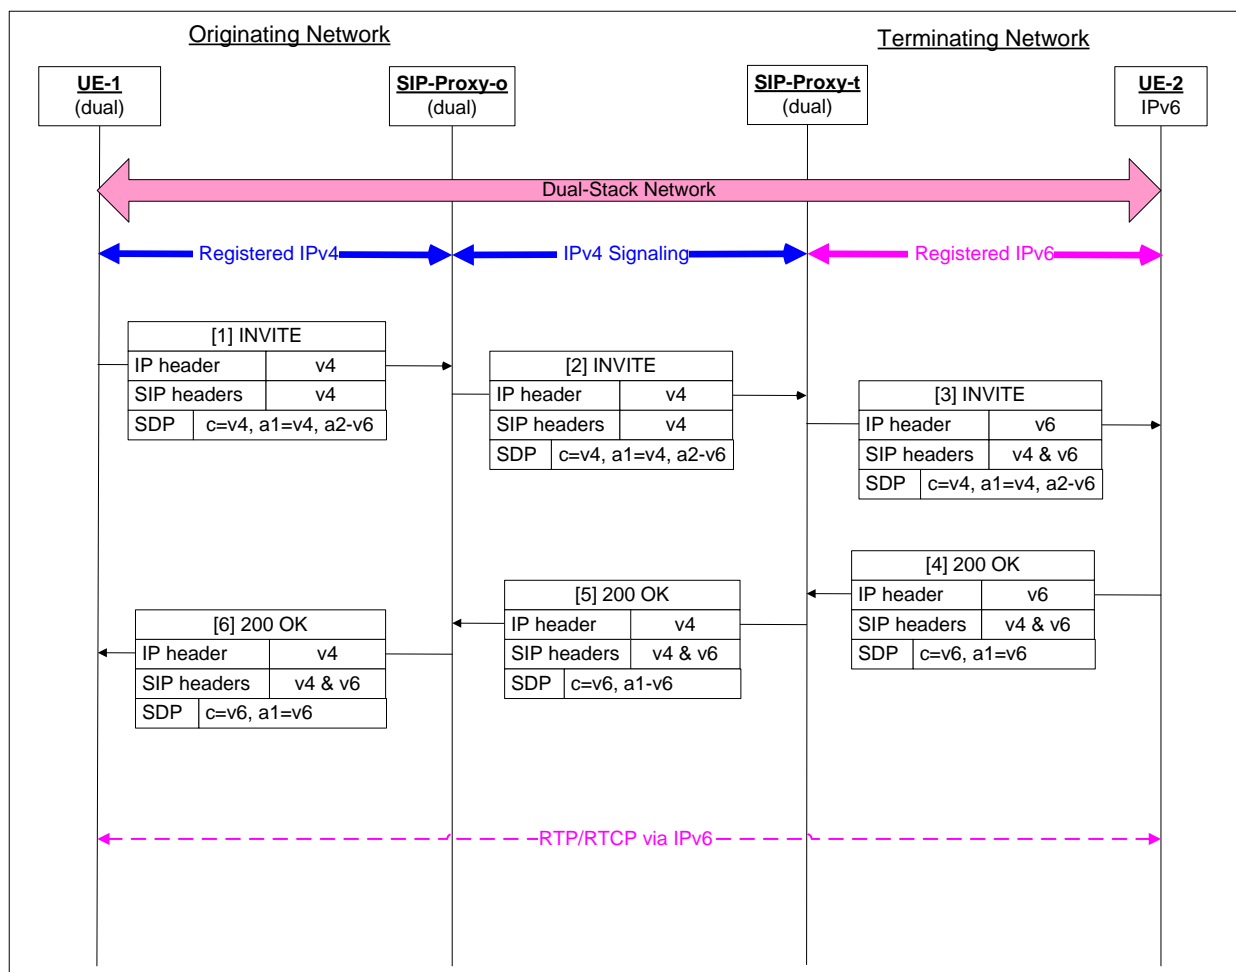


Figure 15 - Dual-Stack UE Registered Using IPv4 Calls Single-Stack IPv6 UE

Initial state:

- UE-1 is dual-stack registered via IPv4 while UE-2 is single stack IPv6.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv4.

Message Sequence:

UE-1 includes uses ICE-lite attributes to advertise both of its IP addresses in the SDP offer. UE-2 also supports ICE-lite, and so responds with an SDP answer containing its IPv6 address as the only candidate, and the session is established over IPv6.

8.3.1.5 Calls Between Mixed IP Version Single-Stack UEs

Figure 16 shows the session establishment signaling-flow between two single-stack endpoints – one IPv4 and one IPv6 – over a dual-stack network.

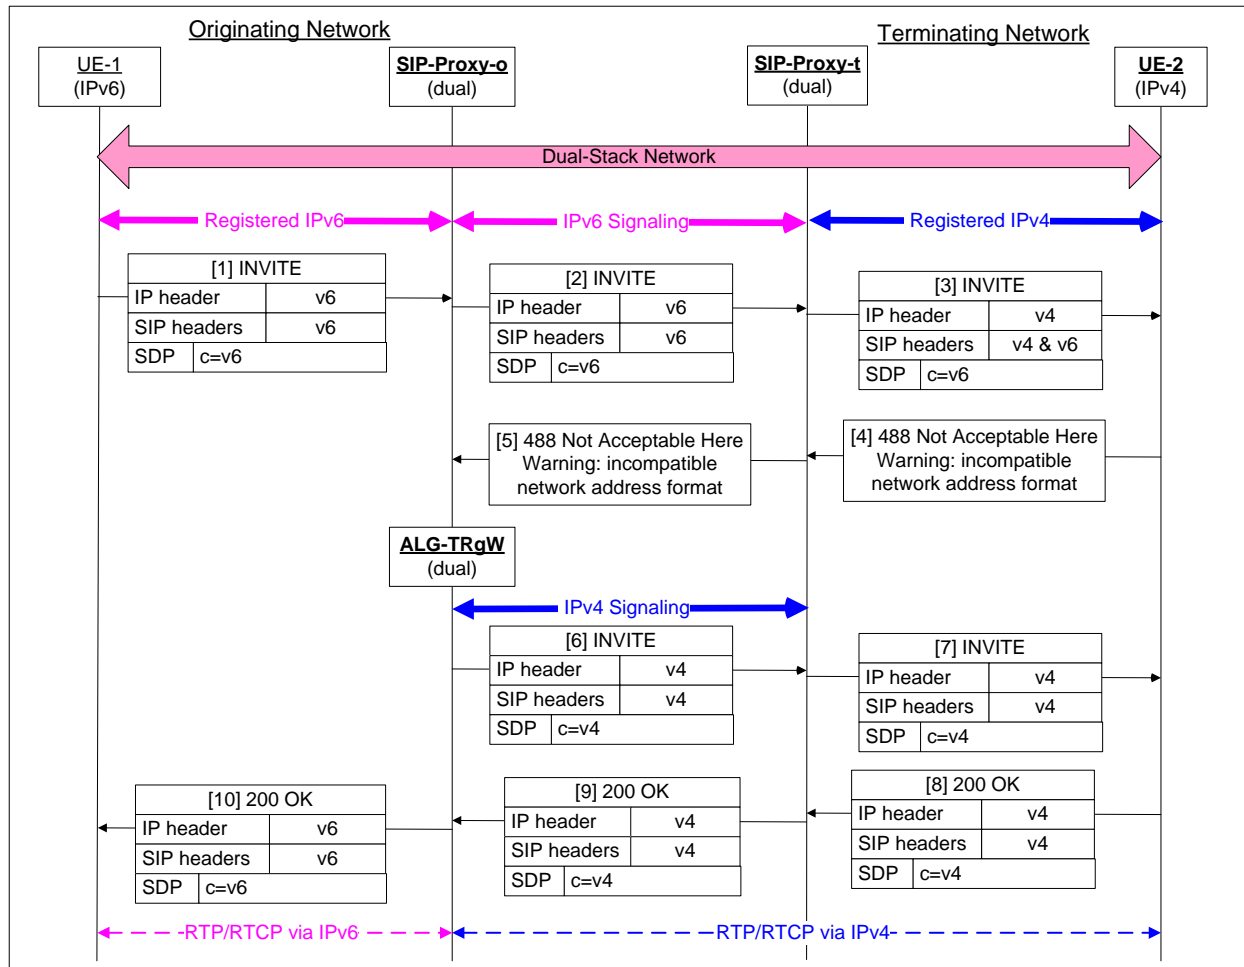


Figure 16 - Single-Stack IPv6 UE Calls Single-Stack IPv4 UE

Initial state:

- UE-1 is single-stack IPv6 while UE-2 is single stack IPv4.
- SIP routing database dictates that SIP-Proxy-o and SIP-Proxy-t signal over IPv6., while the ALG-TrGW and SIP-Proxy-t signal over IPv4.

Message Sequence:

UE-1 offers its IPv6 address in request [1], which is rejected by UE-2 with response [4]. On receiving [5], the originating network inserts an ALG-TrGW, and offers an IPv4 media address in [6]. The offer/answer exchange is completed, and the resulting session has two segments; an IPv6 leg on the originating side of the TrGW, and an IPv4 leg on the terminating side of the TrGW.

8.3.2 Establishing Calls over a Mix of Single- and Dual-Stack Networks

It will be impractical to instantaneously update the entire PacketCable 2.0 core network across all cable operators to dual-stack. Operators will adopt different timelines to migrate their networks to dual-stack, with the result that there

will be a transition period where dual-stack and single-stack IPv4-only networks co-exist. This mix of dual and single-stack networks could also exist within an operator's network, where the operator adopts a strategy to update the network to dual-stack in segments.

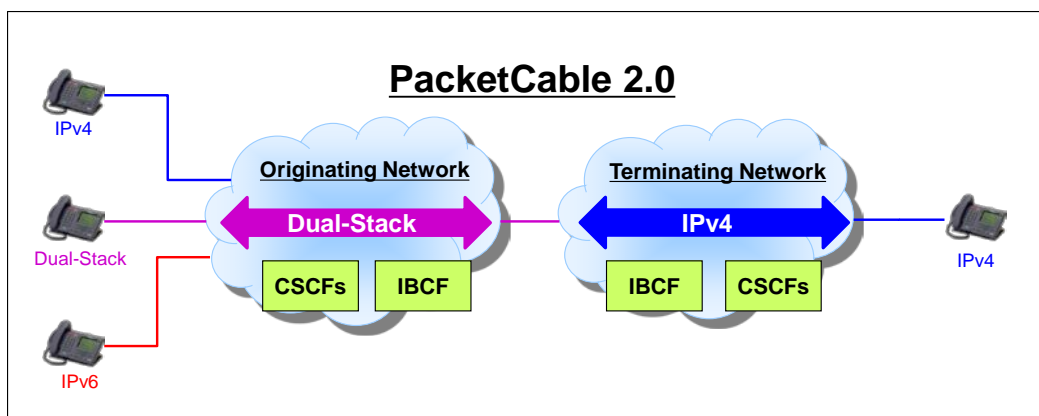


Figure 17 - Mixed Single and Dual-Stack Networks

8.3.2.1 Call from Dual-Stack UE to Single-Stack IPv4 UE

Figure 18 shows the session establishment signaling-flow between a dual-stack endpoint in a dual-stack network, and single-stack IPv4 endpoint in a single-stack network. An ALG is statically inserted to perform IP version interworking.

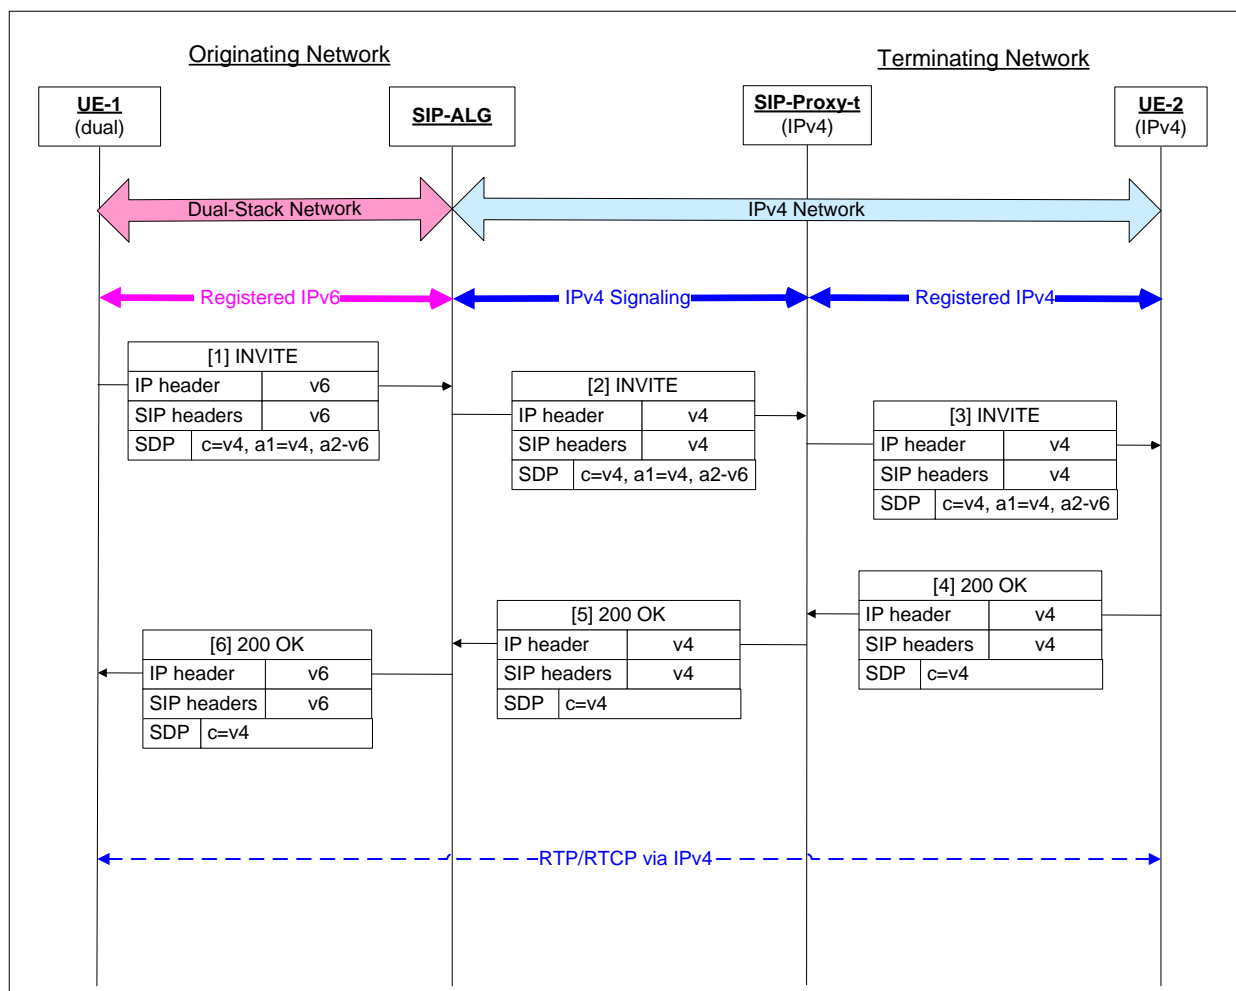


Figure 18 - Dual-Stack UE Calls Single-Stack IPv4 UE over Mixed IPv4/6 Network

Initial state:

- UE-1 is dual-stack registered via IPv6 while UE-2 is single stack IPv4.
- SIP ALG (e.g., IBCF) is statically deployed to perform IP version interworking between originating and terminating networks.

Message Sequence:

UE-1 sends an SDP-offer containing its ICE-lite candidates to UE-2 via [1], [2], [3]. Because the SIP-ALG policy is to perform IP version interworking only (no topology hiding required in this use-case), it does not insert a media-relay at this point. The SIP ALG does do IP version interworking of the IP addresses in the SIP headers, to protect non-compliant endpoints in the terminating network that may not be able to support IPv6 addresses in received SIP header fields. UE-2 returns an SDP-answer containing its IPv4 media address to UE-1 via [4], [5], and [6]. The ALG notices that the two endpoints support compatible media IP versions, and therefore does not insert a media-relay TrGW, but lets the RTP flow end-to-end.

8.3.2.2 Call from Single-Stack IPv6 UE to Single-Stack IPv4 UE

Figure 19 shows the session establishment signaling-flow between a single-stack IPv6 endpoint in a dual-stack network and a single-stack IPv4 endpoint in a single-stack network. An ALG is statically inserted to perform IP version interworking.

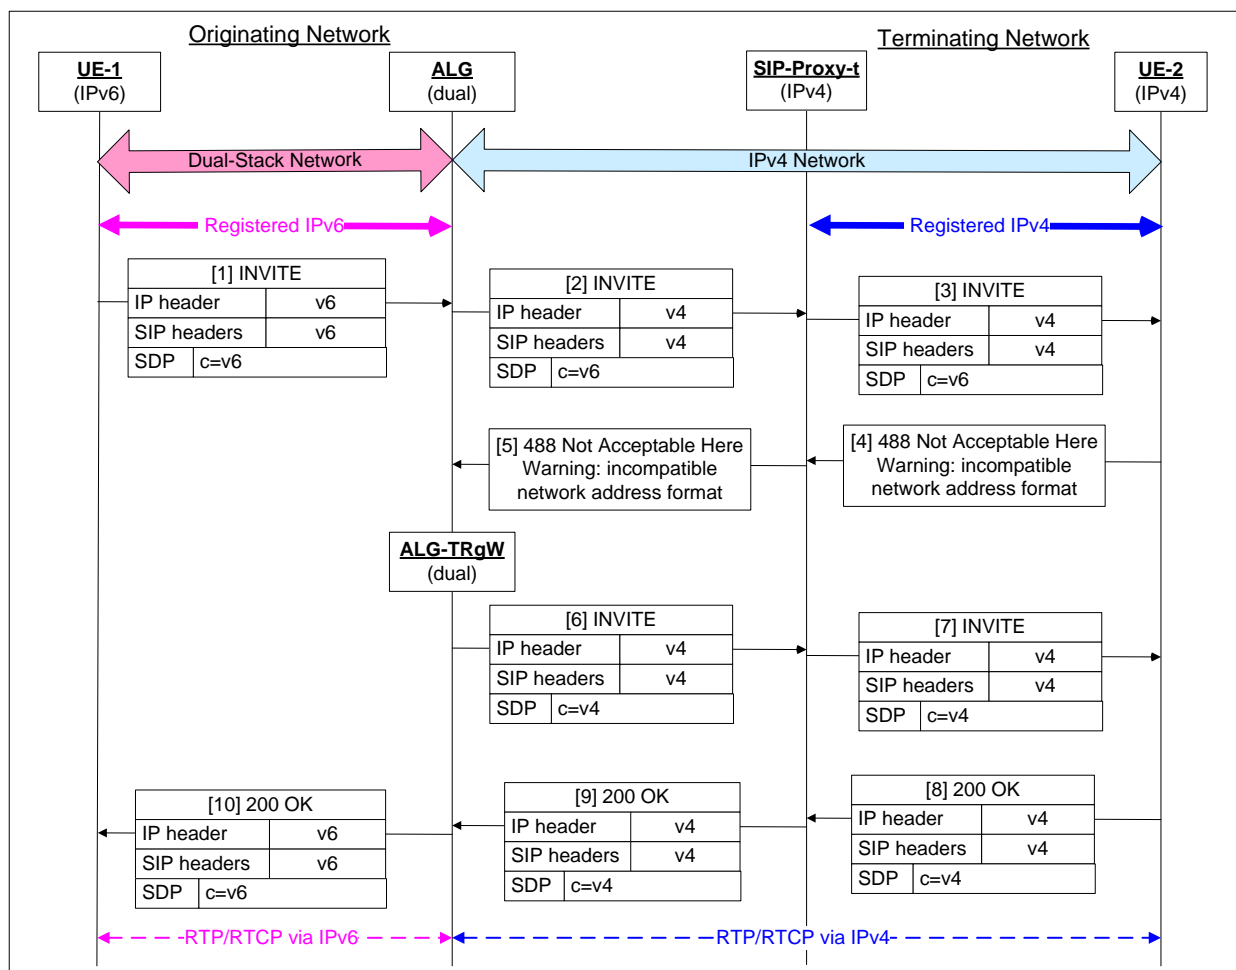


Figure 19 - Single-Stack IPv6 UE Calls Single-Stack IPv4 UE over Mixed Network

Initial state:

- UE-1 is dual-stack registered via IPv6 while UE-2 is single stack IPv4.
- SIP ALG (e.g., IBCF) is statically deployed to perform IP version interworking between originating and terminating networks.

Message Sequence:

UE-1 sends an SDP-offer containing its ICE-lite candidates to UE-2 via [1], [2], [3]. Because the SIP-ALG policy is to perform IP version interworking only (no topology hiding required in this use-case), it does not insert a media-relay. The SIP ALG does do IP version interworking of the IP addresses in the SIP headers, to protect non-compliant endpoints in the terminating network that may not be able to support IPv6 addresses in received SIP header fields. UE-2 returns an error response at [4] indicating media IP version incompatibility. The SIP-ALG inserts a media-relay TrGW to perform IP version interworking of the RTP/RTCP IP header, and the media session is established in two segments – one IPv6 and one IPv4 – interworking at the TrGW.

8.3.3 Interworking Between PacketCable 1.5 and 2.0

PacketCable 1.5 networks will not be upgraded to dual-stack. The session establishment signaling flows will be identical to those between dual-stack and single-stack-IPv4 PacketCable 2.0 networks, as shown in Section 8.3.2.

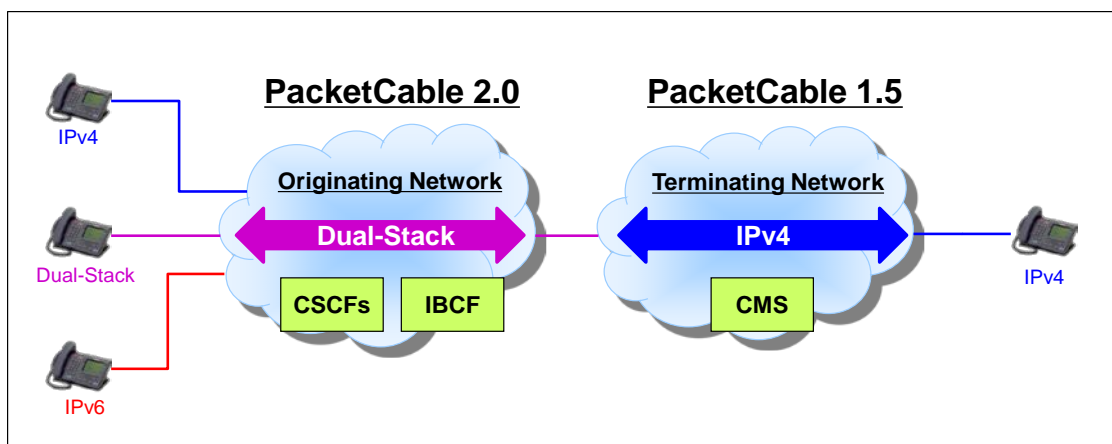


Figure 20 - Interworking Between PacketCable 1.5 and Dual-Stack PacketCable 2.0 Networks

8.3.4 Two Dual-Stack Networks Connected by an IPv4 Transit Network

Originating and Terminating Dual-Stack PacketCable 2.0 Networks may be interconnected using an IPv4 IMS Transit Network. In this case, it is assumed that the IPv4 IMS Transit network provides Ingress/Egress IBCF functions with an associated IPv4 media plane border function. As such, from the perspective of the Originating and Terminating PacketCable 2.0 Networks, session establishment signaling flows will be identical to those between dual-stack and single-stack-IPv4 PacketCable 2.0 networks.

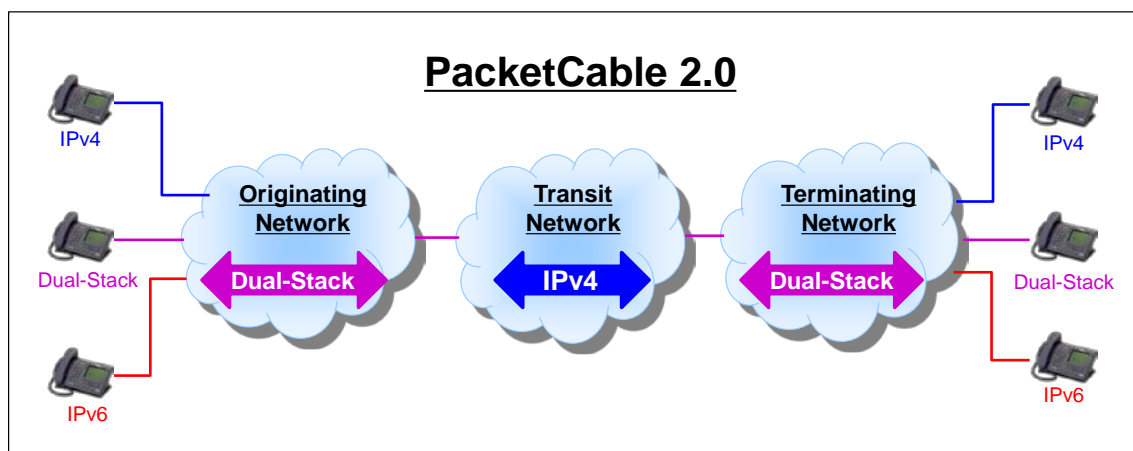


Figure 21 - Dual-Stack Networks Interconnected by Single-Stack Transit Network

Appendix I Acknowledgements

CableLabs wishes to acknowledge the following Dual-Stack focus team members for their contributions to the development of this Technical Report:

Gordon Li (Broadcom)
Eugene Nechamkin (Broadcom)
Jim Shoghli (Broadcom)
Dan Wing (Cisco)
Brian Lindsay (GENBAND)
Dany Sylvain (GENBAND)
Satish Mudugere (Intel)
Jim Stanco (Nokia Siemens)
Geoff Devine (SMC)
Carl Klatsky (Comcast)
Clark Whitten (Cox)
Leonard Mosley (TWC)
John Berg (CableLabs)
Eduardo Cardona (CableLabs)
Chris Grundemann (CableLabs)

David Hancock and the PacketCable Architects
