

Superseded **by a later version of this document**

PacketCable™ 2.0

Quality of Service Specification

PKT-SP-QOS-I01-070925

ISSUED

Notice

This PacketCable specification is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. 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 2007 Cable Television Laboratories, Inc.
All rights reserved.

Document Status Sheet

Document Control Number:	PKT-SP-QOS-I01-070925			
Document Title:	Quality of Service Specification			
Revision History:	I01 - September 25, 2007			
Date:	September 25, 2007			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author Only	CL/Member	CL/Member/Vendor	Public

Key to Document Status Codes

- Work in Progress** An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
- Draft** A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
- Issued** A stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and for certification testing.
- Closed** A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

CableLabs®, DOCSIS®, EuroDOCSIS™, eDOCSIS™, M-CMTS™, PacketCable™, EuroPacketCable™, PCMM™, CableHome®, CableOffice™, OpenCable™, OCAP™, CableCARD™, M-Card™, and DCAS™ are trademarks of Cable Television Laboratories, Inc.

Contents

1	SCOPE	1
1.1	Introduction and Purpose.....	1
1.2	Organization of document.....	1
1.3	Requirements.....	1
2	REFERENCES	3
2.1	Normative References.....	3
2.2	Informative References.....	3
2.3	Reference Acquisition.....	3
3	TERMS AND DEFINITIONS	4
4	ABBREVIATIONS AND ACRONYMS	5
5	OVERVIEW	6
5.1	QOS Architecture Reference Model.....	6
5.2	Functional Components.....	8
5.2.1	<i>P-CSCF</i>	8
5.2.2	<i>PacketCable Application Manager (PAM)</i>	8
5.2.3	<i>Policy Server (PS)</i>	9
5.2.4	<i>Cable Modem Termination System (CMTS)</i>	9
5.2.5	<i>Relationship between PCC Rules and PacketCable Multimedia Gates</i>	9
5.2.6	<i>Relationship between P-CSCF and PAM</i>	9
5.3	Protocol Interfaces and Reference Points.....	10
5.3.1	<i>P-CSCF - PAM Interface Description</i>	10
5.4	Application of QoS Policy within PacketCable.....	10
6	PACKETCABLE SUPPORT OF 3GPP Rx REFERENCE POINT	12
6.1	P-CSCF.....	12
6.1.1	<i>Support for Network Enforced Preconditions</i>	12
6.1.2	<i>Initial Provisioning of Session Information</i>	12
6.1.3	<i>Provisioning of Service Information at the PAM</i>	13
6.2	PacketCable Application Manager.....	13
6.2.1	<i>Initial Provisioning of Session Information</i>	14
6.2.2	<i>Modification of Session Information</i>	14
6.2.3	<i>AF Session Termination</i>	15
6.2.4	<i>Subscription to Notification of Signaling Path Status</i>	15
7	PARAMETER MAPPING	16
7.1	Mapping AA-Requests to Gate-Set.....	16
7.1.1	<i>FlowSpec Traffic Profile</i>	16
7.1.2	<i>Classifiers</i>	20
7.1.3	<i>GateSpec</i>	21
7.1.4	<i>SubscriberID</i>	21
7.1.5	<i>Application Type ID</i>	21
7.2	Mapping Gate-Set-Ack to AA-Answers.....	21
7.2.1	<i>Access Network Charging Identifier</i>	21
	APPENDIX I ACKNOWLEDGEMENTS	22

Figures

Figure 1 - QoS Architecture Reference Model.....7
Figure 2 - One-to-One Relationship between P-CSCF and PAM.....9
Figure 3 - Many-to-One Relationship between P-CSCF and PAM.....9

Tables

Table 1 - QoS Reference Points7

1 SCOPE

1.1 Introduction and Purpose

The purpose of this specification is to define the requirements and functionality needed to support Quality of Service (QoS) within the PacketCable 2.0 Architecture. PacketCable 2.0 is a CableLabs specification effort designed to extend cable's real-time IP communication service architecture and to accelerate the convergence of voice, video, data, and mobility technologies which makes use of the rich set of QoS capabilities provided by PacketCable Multimedia.

The primary focus of this document is to define how the QoS capabilities of PacketCable Multimedia are integrated with the Policy and Charging Control (PCC) architecture defined by the 3rd Generation Partnership Project (3GPP) for its IP Multimedia Subsystem (IMS) which forms the basis of the PacketCable 2.0 control plane architecture. In addition to describing how the PCC architecture is realized in a PacketCable 2.0 network and providing an overview of network elements which make up the architecture, this specification defines detailed requirements for passing session level information by the Proxy Call Session Control Function (P-CSCF) and the use of this information by the PacketCable Application Manager (PAM) to provide QoS in the DOCSIS® access network.

The architecture and requirements defined in this specification allow operators to take advantage of the rich QoS capabilities of DOCSIS, while reaping the economies of scale and scope of deploying an access-agnostic services infrastructure which does not require cable specific development to enable QoS.

1.2 Organization of document

Section 5 provides an architectural overview of the elements involved in the QoS architecture.

Section 6 provides normative requirements for both the P-CSCF and the PAM.

Section 7 provides normative requirements for the mapping between the Rx and PacketCable Multimedia interfaces.

1.3 Requirements

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

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

"MAY"

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

2 REFERENCES

2.1 Normative References

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

- [PCMM] PacketCable Multimedia Specification, PKT-SP-MM-I03-051221, December 21, 2005, Cable Television Laboratories, Inc.
- [TS 29.214] 3GPP TS 29.214, Policy and Charging Control over Rx reference point, Release 7, V7.2.0, September 2007.
- [TS 29.213] 3GPP TS 29.213, Policy and Charging Control signalling flows and QoS parameter mapping, Release 7, V7.2.0, September 2007.

2.2 Informative References

This specification uses the following informative references.

- [CPD] PacketCable Control Point Discovery Specification, PKT-SP-CPD-I03-070925, September 25, 2007, Cable Television Laboratories, Inc.
- [CODEC-MEDIA] PacketCable Specification, Codec and Media Specification, PKT-SP-CODEC-MEDIA-I03-070925, September 25, 2007, Cable Television Laboratories, Inc.
- [TS 23.203] 3GPP TS 23.203, Policy and charging control architecture, Release 7, V7.3.0, June 2007.
- [PKT 24.229] PacketCable SIP and SDP Stage 3 Specification 3GPP TS 24.229, PKT-SP-24.229-I03-070925, September 25, 2007, Cable Television Laboratories, Inc.
- [TS 29.212] 3GPP TS 29.212, Policy and Charging Control over Gx reference point, Release 7, V7.2.0, September 2007.
- [ICE] IETF Internet Draft, Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols, draft-ietf-mmusic-ice-18.txt, September 2007, work in progress.
- [NAT] PacketCable NAT and Firewall Traversal Technical Report, PKT-TR-NFT-V03-070925, September 25, 2007, Cable Television Laboratories, Inc.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; <http://www.cablelabs.com>
- Third Generation Partnership Project (3GPP), Internet: <http://www.3gpp.org>
- Internet Engineering Task Force (IETF), Internet: <http://www.ietf.org/>
Note: Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time.
The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

3 TERMS AND DEFINITIONS

This specification uses the following terms:

DOCSIS	A set of interface specifications for transmitting data over cable television systems in a standard fashion.
Network Enforced Preconditions	A method of ensuring that necessary QoS resources are available before alerting the called user by withholding SIP messages within the network until the results of necessary policy interactions are known. This capability is independent from the use by the UE of the similarly named QoS Preconditions capability defined by RFCs 3312 and 3313.
PacketCable Application Manager (PAM)	A specialized Application Manager defined in PacketCable, primarily responsible for determining QoS resources needed for a session.
Policy Server	A system that primarily acts as an intermediary between Application Manager(s) and CMTS(s). It applies network policies to Application Manager requests and proxies messages between the Application Manager and CMTS.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

AF	Application Function
AMID	Application Manager Identifier
BCID	Billing Correlation ID
CM	DOCSIS-compliant cable modem
CMTS	Cable Modem Termination System
CSCF	Call Session Control Function
IP-CAN	IP Connectivity Access Network
DOCSIS®	Data Over Cable Service Interface Specifications
DSCP	Differentiated Services Code Point
NAT	Network Address Translation
PAM	PacketCable Application Manager
PCC	Policy and Charging Control
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rules Function
P-CSCF	Proxy CSCF
PS	Policy Server
SDP	Session Description Protocol
SIP	Session Initiation Protocol
SPR	Subscription Profile Repository
UE	User Equipment

5 OVERVIEW

PacketCable is a CableLabs specification effort designed to support the convergence of voice, video, data, and mobility technologies. This release of PacketCable 2.0 is based on the IP Multimedia Subsystem (IMS) as developed by the 3rd Generation Partnership Project (3GPP) while making use of the rich set of QoS capabilities provided by PacketCable Multimedia.

PacketCable Multimedia defines an IP-based platform for delivering QoS-enhanced multimedia services over DOCSIS 1.1 or greater access networks (for the remainder of this document, references to DOCSIS assume DOCSIS 1.1 or greater).

The QoS and event messaging capabilities of PacketCable Multimedia are integrated into the IMS based PacketCable architecture through the Policy and Charging Control (PCC) capability introduced in 3GPP release 7. PCC represents the evolution of the separate policy and charging control mechanisms defined in earlier 3GPP releases while incorporating additional requirements, such as support for policy push, indication of session priority, more detailed session description information, and support for separate reservation and commitment of resources which allow PCC to support various access technologies.

5.1 QOS Architecture Reference Model

The PacketCable QoS architecture reference model, illustrated in Figure 1, uses the generic PacketCable Multimedia components including the Policy Server, CMTS, and Cable Modem. The PacketCable Application Manager (PAM) is a specialized Application Manager that receives session-level QoS requests via Rx from the P-CSCF and creates and manages the PacketCable Multimedia Gates for each flow in the session using the PacketCable Multimedia pkt-mm-3 interface. The PacketCable Application Manager interacts with the P-CSCF via Rx for sessions involving UEs in the PacketCable Multimedia network just as the PCRF does for sessions involving UEs in the GPRS network.

Figure 1 illustrates how the PacketCable Multimedia QoS architecture, used for UEs accessing services via a DOCSIS network, integrates into the 3GPP Policy and Charging Control architecture.

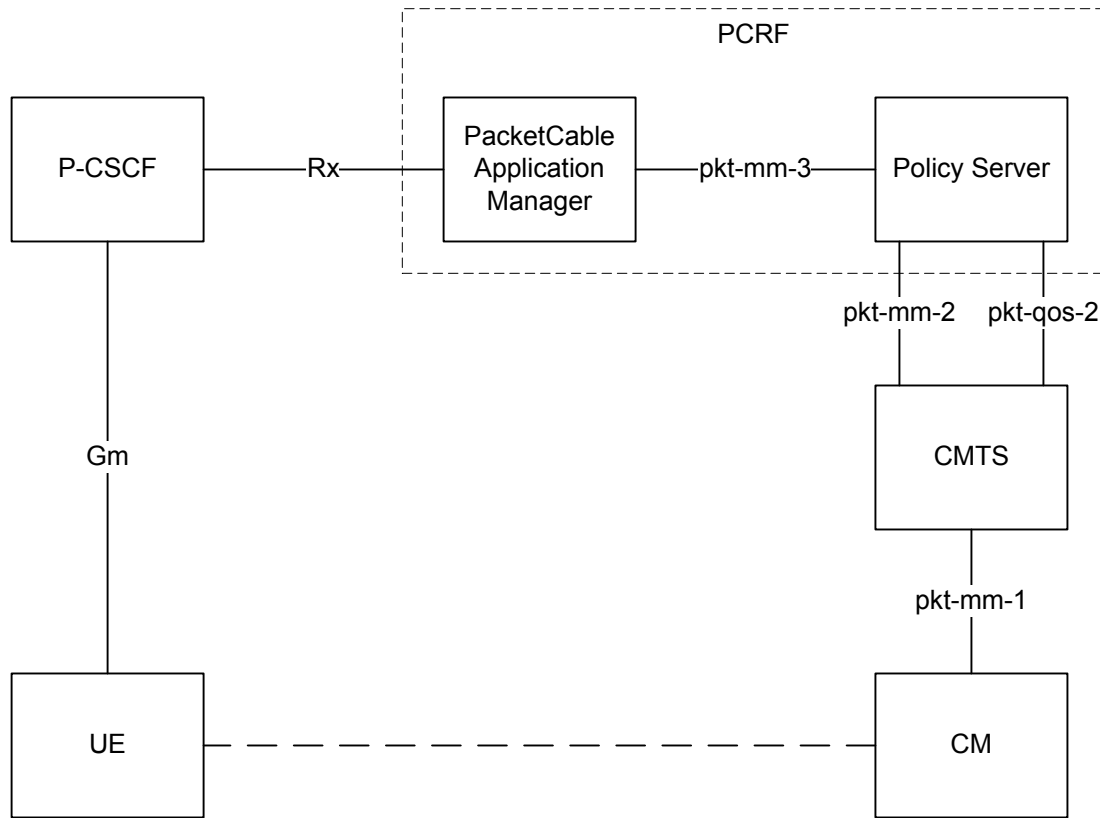


Figure 1 - QoS Architecture Reference Model

Reference points pkt-mm-1 through pkt-mm-3 are defined by PacketCable Multimedia. Reference point pkt-qos-2 is defined by PacketCable 2.0. Reference points Gm and Rx are defined by 3GPP. Table 1 describes these interfaces briefly.

Table 1 - QoS Reference Points

Reference Point	PacketCable Network Elements	Reference Point Description
pkt-mm-1	CMTS - CM	The CMTS uses DOCSIS-defined DSX signaling to instruct the CM to setup, teardown, or change a DOCSIS service flow in order to satisfy a QoS request.
pkt-mm-2	Policy Server - CMTS	The interface supports proxy QoS requests on behalf of a UE. This interface is fundamental to the policy-management framework. It controls policy decisions, which are pushed by the Policy Server (PS) onto the CMTS, and is defined by [PCMM]. In some scenarios, this interface is also used to inform the PS of changes of the status of QoS resources.
pkt-mm-3	PAM - PS	This interface allows the PacketCable Application Manager (PAM), a specialized Application Manager defined in PacketCable, to request that the PS install a policy decision on the CMTS on behalf of the UE, and is defined by [PCMM]. This interface may also be used to inform the PAM of changes in the status of QoS resources.

Reference Point	PacketCable Network Elements	Reference Point Description
Gm	UE - P-CSCF	Allows the UE to communicate with the P-CSCF for registration and session control. This reference point is SIP-based and is described in [PKT 24.229].
Rx	P-CSCF - PAM	The Rx interface is used for session-based policy set-up information exchange between the P-CSCF and the PacketCable Application Manager. Note that this is the same interface that is used between the P-CSCF and the Policy and Charging Rules Function (PCRF) in the case of GPRS access. This interface may also be used to inform the P-CSCF of changes in the status of QoS resources.
pkt-qos-2	Policy Server - CMTS	The Policy Server uses the Control Point Discovery Protocol to determine the serving CMTS in the network for a given UE. This reference point is based on PacketCable specification [CPD].

5.2 Functional Components

In this section, additional detail is provided on the P-CSCF, PacketCable Application Manager, Policy Server, and CMTS and their roles as related to the QoS architecture.

5.2.1 P-CSCF

The P-CSCF fulfills the role of the Application Function (AF) in the 3GPP PCC architecture [TS 23.203].

In addition to its role in providing the UE connectivity to the PacketCable network, the P-CSCF is also responsible for reserving, committing, and releasing QoS resources for a given session. It is important to note that the P-CSCF does not actually determine the QoS resources necessary for the session; rather, it simply proxies the session description information to the PacketCable Application Manager and indicates whether to reserve or commit the resources for the session. While the architecture supports a two-phase commit operation (reserve followed by a commit), there are no requirements on the P-CSCF to follow this approach. A single-phase commit (reserve and commit resources in a single request) may be used.

Once the session has ended, the P-CSCF informs the PAM, which releases the resources allocated to the session. Similarly, the PAM may be requested to inform the PCSCF if resources allocated for the session are lost during the session. The P-CSCF may then initiate procedures to tear down the session.

5.2.2 PacketCable Application Manager (PAM)

The PAM, in conjunction with the Policy Server (PS), fulfills the role of the Policy and Charging Rules Function (PCRF) in the 3GPP IMS PCC architecture [TS 23.203].

The PacketCable Application Manager is primarily responsible for determining the QoS resources needed for the session, based on the received session descriptors from P-CSCF, and managing the QoS resources allocated for a session.

Determining the QoS resources for a session involves interpreting the session information and calculating how much bandwidth is required, determining the correct PacketCable Multimedia traffic profile, and populating the traffic classifiers. This also involves determining the number of flows necessary for the session (e.g., voice only vs. voice and video session) and managing the association of the flows to the session.

5.2.3 Policy Server (PS)

The PacketCable 2.0 QoS architecture makes use of the PacketCable Multimedia Policy Server (PS) defined in [PCMM].

5.2.4 Cable Modem Termination System (CMTS)

The CMTS fulfills the role of Policy and Charging Enforcement Function (PCEF) in the 3GPP PCC architecture [TS 23.203]. As described in the following section, enforcement at the CMTS is done on the basis of PacketCable Multimedia Gates as defined in [PCMM] rather than PCC rules.

5.2.5 Relationship between PCC Rules and PacketCable Multimedia Gates

The PacketCable 2.0 QoS architecture makes use of PacketCable Multimedia Gates to manage QoS and, through event messaging, charging for IP flows associated with a particular session. While these Gates are analogous and provide similar functionality to the PCC Rules exchanged over the Gx reference point in the PCC architecture as used with other access technologies, there are some basic differences. PacketCable Multimedia Gates do not follow the same format as PCC rules. PacketCable Multimedia Gates are unidirectional so that separate Gates are required for the upstream and downstream flows associated with a single bidirectional media session. In addition, DOCSIS does not support the concept of predefined PCC rules.

5.2.6 Relationship between P-CSCF and PAM

The PacketCable Multimedia architecture provides a well-understood relationship between the PAM and PS. The relationship between the P-CSCF and PAM is described here. The QoS architecture was not developed with any pre-conceived relationship between the two network elements. While the choice of how to deploy P-CSCFs and PAMs and their associated cardinality is mainly a deployment decision, the following figures represent what are believed to be the most popular deployment scenarios.




Figure 2 - One-to-One Relationship between P-CSCF and PAM




Figure 3 - Many-to-One Relationship between P-CSCF and PAM

Figure 2 illustrates a one-to-one relationship between the P-CSCF and the PAM. Such a deployment scenario, while extremely simple to manage, may not be the most efficient use of resources. Figure 3 illustrates a many-to-one relationship between the P-CSCF and the PAM. This scenario maintains the simplicity of a one-to-one relationship (since the P-CSCF does not have to determine which PAM to send its request to), but may more efficiently utilize PAM resources.

Another possible scenario not shown is a one-to-many relationship (or many-to-many). In this case, the P-CSCF may communicate with multiple PAMs. While such a scenario is supported, no guidance is provided on how the P-CSCF determines which PAM to send its requests to. Such a scenario may become necessary as the network evolves and different application managers are invoked for different access networks, or in certain roaming cases.

5.3 Protocol Interfaces and Reference Points

This specification has identified several interfaces, or reference points, in the PacketCable QoS architecture. The majority of these reference points are existing interfaces defined by PacketCable Multimedia. An overview of the P-CSCF to PAM protocol interface is provided within this section.

5.3.1 P-CSCF - PAM Interface Description

The P-CSCF to PAM interface is an instance of the DIAMETER based Rx interface defined in 3GPP release 7. This interface allows the P-CSCF to provide session information to the PAM and to receive from the PAM information about the state of QoS resources associated with that session. The PAM manages the determination, reservation, commitment and release of necessary Quality of Service resources to support the session within a PacketCable Multimedia-enabled DOCSIS network via pkt-mm-3.

5.4 Application of QoS Policy within PacketCable

The term policy control has often been used to describe the process by which new dynamic services flows or bearers are created in the access network at the request of an application. This makes sense, since establishing a new service flow in the access network involves the installation of a new dynamic policy in the policy enforcement point. This dynamic policy determines the treatment of the packets that make up the new service flow within the access network throughout the duration of the session.

The focus of this section is on higher levels of policy, which may affect the disposition of a user request as it is processed within the network. Such policies could be implemented at several levels in the network in order to further the business needs of the network operators. Levels where policy can be applied include:

- **Application Level:** Applications may employ policy to constrain use of an application based on subscription or other information.
- **Signaling Network Level:** Network or subscription-based restrictions on the use of certain media parameters in an SDP offer may be enforced in a PacketCable network by the P-CSCF or S-CSCF, respectively, by sending a negative response to the SIP message as described in [PKT 24.229] clauses 6.2 and 6.3.
- **Resource Control Level:** Each of the network elements in the bearer network (PacketCable Application Manager, Policy Server and CMTS) serve unique roles as they relate to policy control. A more detailed discussion of the roles for each network element follows:
 - The PacketCable Application Manager is the entry point from the IMS network into the cable access network QoS system. The PAM is in a position to apply policy, taking into consideration the service information provided by the P-CSCF and potentially subscriber-based information.
 - The Policy Server may receive messages from multiple application managers, including but not limited to PacketCable Application Managers. Policies applied at the PS can optimize the use of access network resources between multiple applications and traffic types.

- The CMTS is responsible for admission control and may have local policies that control the allocation of resources among various types of traffic based on PacketCable Multimedia session class and possibly on the authorization model used, such as PacketCable DQoS and PacketCable Multimedia.

In some cases similar policy decisions could be made at more than one level in the network. The choice at which level to implement a given policy will be based on such criteria as:

- Access to required information.
- Performance impacts of implementing policy at that level.
- Ease of implementing policy at that level.

Ultimately policy will be implemented at the level in the network that meets the business needs of the cable operator in the most optimal way.

6 PACKETCABLE SUPPORT OF 3GPP Rx REFERENCE POINT

The Rx reference point is used to exchange application level session information between the P-CSCF and the PacketCable Application Manager (PAM). The PAM uses this application level session information to create PacketCable Multimedia Gates on the CMTS which is providing access network resources for the session. This section documents the relevant P-CSCF and PAM requirements as defined in 3GPP [TS 29.214] and [TS 29.213], as well as any deltas to the 3GPP specifications.

6.1 P-CSCF

The P-CSCF implements the Application Function (AF) as defined in [TS 23.203] and [TS 29.214]. When acting as an AF, the P-CSCF is responsible for communicating session resource needs to the PAM via the Rx interface.

The P-CSCF MUST implement the Application Function (AF) requirements as defined in [TS 29.214]. The P-CSCF MUST implement the P-CSCF requirements as defined in Annex A of [TS 29.214]. The P-CSCF MUST implement the QoS parameter mapping functions at the AF as defined in clause 6.2 of [TS 29.213]. In addition to the above referenced 3GPP requirements, the P-CSCF MUST include the Codec-Data AVP in all AA-Requests which are related to sessions which utilize the Session Description Protocol (SDP). This is necessary to ensure the PAM has enough information about the session to efficiently allocate access network resources.

Finally, the P-CSCF MAY implement the call flows as defined in Annex B of [TS 29.213]. Please see 6.1.1 for details on the interactions between network enforced preconditions and these call flows.

6.1.1 Support for Network Enforced Preconditions

Network enforced preconditions is a method of preventing the alerting of a called user until necessary network resources have been reserved. This is done by withholding SIP messages that contain an SDP offer or SDP answer until the results of any necessary policy interaction is known. QoS resources are reserved based on information in the SDP offer. After the QoS reservation is successfully completed, the SIP message can be forwarded to the next Hop device (UE or upstream Network Element).

Implementation of the call flows in Annex B of [TS 29.213] do not allow for the use of network enforced preconditions, since the session is allowed to progress to an alerting state before resources are reserved. In order to allow for network enforced preconditions the P-CSCF MUST provide session information to the PAM upon receipt of an SDP Offer so that the PAM can reserve access network resources. The P-CSCF MUST wait for a success indication from the PAM before allowing the session to proceed. If the P-CSCF receives an indication other than success from the PAM, the P-CSCF SHOULD reject the session attempt. Upon receipt of an SDP Answer, the P-CSCF MUST update the previously provided session information to allow the PAM to release any extra reserved resources that will not be used by the session and install a more descriptive packet classifier.

6.1.2 Initial Provisioning of Session Information

3GPP [TS 29.214] clause 4.4.1 specifies that the AF shall provide the UE's IP address using either Framed-IP-Address AVP or Framed-IPv6-Prefix AVP. No guidance on how the P-CSCF determines the UE's IP address is provided at this time. Normally the UE's IP address can be either taken from the address in the c= line of the UE provided SDP, or as the source IP Address of the UE's signaling exchange with the P-CSCF. However, given that NATs may be employed between the UE and the access network, neither of these addresses are guaranteed to be correct (the address in the c= line of the SDP may be that of a relay server while the source IP Address for signaling may be different than the IP Address used for Media). In order to obtain a valid, network assigned, IP address which allows the PS to identify the proper CMTS serving the UE, the P-CSCF MUST perform as follows:

1. Inspect the UE provided SDP for a=candidate attributes. If such attributes are present, then the UE has employed the ICE procedures [ICE] and the following procedures apply for determining the proper UE provided IP address for either the Framed-IP-Address AVP or Framed-IPv6-Prefix AVP:
 1. Inspect each candidate entry for a candidate type of server reflexive (“srflx”). If such an entry is present, then the UE IP Address MUST be that as provided as the server reflexive candidate.
 2. If there are no candidate entries of type server reflexive, then inspect the each entry for a candidate type of “relay”. If such an entry is found, then the UE IP Address MUST be that as provided by the rel-addr property of this candidate.
 3. If there are no candidates of type server reflexive or relay, then the UE IP Address MUST be that as provided by the address in c= line.
2. If no candidate attributes are present in the UE provided SDP, then the UE has not employed the ICE procedures and the P-CSCF MUST use the IP address contained in the c= line for either the Framed-IP-Address AVP or Framed-IPv6-Prefix AVP.

6.1.3 Provisioning of Service Information at the PAM

3GPP [TS 29.214] Annex A.1 specifies how the P-CSCF shall derive the Flow-Description AVP. However, this text does not consider the case where a NAT is present between the network and the UE and the UE has utilized the ICE procedures. When a UE utilizes the ICE procedures, it is possible that the address and port advertised in the 'c=' and 'm=' lines of the SDP is actually that of a STUN Relay server and not that of the UE or the WAN interface of the NAT. See the NAT and Firewall Traversal Technical Report [NAT] for further details on this issue.

In the case where STUN Relay and ICE are used for NAT traversal, the P-CSCF MUST inspect the uplink SDP for "a=candidate" attributes and compare the candidate address with that contained in the "c=" line. If a match is found, the P-CSCF then looks at the candidate type. If the candidate type is "relay" then the address in the "c=" line is that of a STUN Relay server. In this case, the P-CSCF MUST derive the Flow-Description AVP within the service information from the SDP candidate type of "relay" as follows:

Uplink Filter

- Destination Address and Port: If the P-CSCF knows the destination address and port of the STUN Relay allocation that the UE is sending media to, it SHOULD use that information. If the P-CSCF does not know this address and port, it MUST wildcard the uplink destination address and port.
- Source Address and Port: The P-CSCF MUST populate the uplink source address with the "rel-addr" address and the uplink source port with the “rel-port” port contained within the “a=candidate” attribute.

Downlink Filter

- Destination Address and Port: The P-CSCF MUST populate the downlink destination address with the "rel-addr" address and the downlink destination port with the "rel-port" port contained within the "a=candidate" attribute.
- Source Address and Port: If the P-CSCF knows the source address and port of the STUN Relay allocation that the UE is receiving media from, it SHOULD use that information. If the P-CSCF does not know this address and port, it MUST wildcard the downlink source address and port.

6.2 PacketCable Application Manager

The PAM is defined by the PacketCable QoS architecture as the entry point to the PacketCable Multimedia architecture. As such, the PAM is responsible for receiving session level resource requests via the Rx interface and translating to PacketCable Multimedia specific requests.

The PAM MUST support the PCRF requirements as defined in [TS 29.214], including Annex A, except as defined in the following sub-sections. As noted in Section 5.2.5, PCC rules are analogous to PacketCable Multimedia Gates. As such, references to PCC rules in [TS 29.214] should be read as PacketCable Multimedia Gates.

6.2.1 Initial Provisioning of Session Information

3GPP [TS 29.214] clause 4.4.1 specifies that once the PCRF determines it needs to provision a PCC rule to the PCEF, the provisioning of the PCC rules to the PCEF be carried out as specified in [TS 29.212] (Gx reference point). In a PacketCable network, the pkt-mm-2 reference point replaces the Gx reference point. Therefore, upon receipt of an AA-Request from the P-CSCF that requires the establishment of a PacketCable Multimedia Gate, the PAM MUST send PacketCable Multimedia Gate-Set requests as defined in [PCMM]. The PAM MUST populate the PacketCable Multimedia Gate-Set Objects as defined in 7.1.

In addition to the procedures defined in clause 4.4.1 of [TS 29.214], the following requirements are also imposed on the PAM:

The PAM SHOULD wait until a response (Gate-Set-Ack, or Gate-Set-Error) is received for all Gate-Set requests before sending an AA-Answer. The decision to wait for all responses before providing an AA-Answer is a matter of operator policy and is necessary if network preconditions are being used. If the PAM chooses to wait until all responses are received, upon receipt of the last outstanding response to a Gate-Set request for a given AA-Request, the PAM MUST follow the following rules:

- If the PAM receives a Gate-Set-Error response to one or more Gate-Set requests, the PAM MUST consider the session modification request failed and return an AA-Answer with an Experimental-Result-Code of 5063 (REQUESTED_SERVICE_NOT_AUTHORIZED). The PAM MUST consider the AA-Request atomic and release any resources established for the session.
- If the PAM is able to successfully execute all of the AA-Request, the PAM MUST return an AA-Answer with a DIAMETER result code of DIAMETER_SUCCESS. In addition, the PAM MUST maintain an association of the DIAMETER session ID and the assigned GateID(s). The PAM SHOULD include the IP-CAN type AVP, as defined in [TS 29.212] in the AA-Answer. If the AA-Request results in DOCSIS resources being created and a Billing Correlation ID being returned via the Access-Network-Charging-Identifier populated as defined in Section 7.2.1, the PAM MUST set the value of the IP-CAN type to DOCSIS (1). Further, the PAM MUST store any relevant information received in the AA-Request such that it can generate any of the messages defined by [TS 29.214], or restore the session state should a subsequent AA-Request to modify the session be unsuccessful.

6.2.2 Modification of Session Information

3GPP [TS 29.214] clause 4.4.2 specifies that the PCRF modify PCC rules as specified in [TS 29.212] (Gx reference point). In a PacketCable network, the pkt-mm-2 reference point replaces the Gx reference point. Therefore, upon receipt of an AA-Request from the P-CSCF which references an existing DIAMETER session ID, the PAM MUST determine whether the request modifies an existing set of Gates, creates a new set of Gates, or removes an existing set of Gates. Once the determination is made, the PAM MUST send PacketCable Gate-Set or Gate-Delete requests as defined in [PCMM]. The PAM MUST populate the PacketCable Multimedia Gate-Set Objects as defined in 7.1.

The PAM SHOULD wait until a response (Gate-Set-Ack, Gate-Set-Error, Gate-Delete-Ack, or Gate-Delete-Error) is received for all Gate-Set or Gate-Delete requests before sending an AA-Answer. The decision to wait for all responses before providing an AA-Answer is a matter of operator policy and is necessary if network preconditions are being used. If the PAM chooses to wait until all responses are received, upon receipt of the last outstanding response to a Gate-Set or Gate-Delete for a given AA-Request, the PAM MUST follow the following rules:

- If the PAM receives a Gate-Set-Error response to one or more Gate-Set requests, the PAM MUST consider the session modification request failed and return an AA-Answer with an Experimental-Result-Code of 5063 (REQUESTED_SERVICE_NOT_AUTHORIZED). The PAM MUST consider the AA-Request

atomic and restore the session to its original state prior to processing the AA-Request modifying existing session resources.

- If the PAM receives a Gate-Delete-Error response to one or more Gate-Delete requests, the PAM MUST consider the Gate deletion successful in determining whether the overall AA-Request is successful based on responses to any Gate-Set requests which may be related with this AA-Request.
- If the PAM is able to successfully execute all of the AA-Request, the PAM MUST return an AA-Answer with a DIAMETER result code of DIAMETER_SUCCESS. In addition, the PAM MUST maintain an association of the DIAMETER session ID and the assigned GateID(s). Further, the PAM MUST store any relevant information received in the AA-Request such that it can generate any of the messages defined by [TS 29.214], or restore the session state should a subsequent AA-Request to modify the session be unsuccessful.

6.2.2.1 Maintenance of PacketCable Multimedia Reserved Timer for Held Media Flows

In some circumstances, for example when a party in a voice session established through residential SIP telephony is placed on hold, a session modification may result in a media flow which was formerly active being made inactive. Such a modification can result in a PacketCable Multimedia gate being transitioned from the committed state to the reserved state. This transition will cause the PacketCable Multimedia Reserved timer (T2) to become active. If no further session modification is performed and no action is taken by the PAM, before timer T2 expires, the gate will be deleted by the CMTS. In order to reduce the likelihood of inadvertent loss of PacketCable Multimedia gates for held session the following behavior is desired.

Upon receipt of an AA-Request which modifies an existing session the PAM may send one or more Gate-Set requests to modify the gates associated with the DIAMETER session. For each of the affected gates, if the resultant state of the gate is Reserved, the PAM MUST periodically re-send the Gate-Set request with a period of less than the value specified for Timer T2. In order to prevent the endless refresh of a reserved gate in a failure condition such as the loss of the P-CSCF, the PAM SHOULD limit the number of times that a reserved gate may be refreshed before letting the T2 Timer expire.

6.2.3 AF Session Termination

3GPP [TS 29.214] clause 4.4.4 specifies that the PCRF utilize the procedures in [TS 29.212] (Gx reference point) to remove PCC rules. In a PacketCable network, the pkt-mm-2 reference point replaces the Gx reference point. Therefore, upon receipt of a ST-Request, the PAM MUST send PacketCable Gate-Delete requests for each Gate associated with the DIAMETER session ID as defined in [PCMM].

6.2.4 Subscription to Notification of Signaling Path Status

The PacketCable Multimedia architecture does not support the subscription to Signaling Path Status. On receipt of an AA-Request as defined in [TS 29.214] clause 4.4.5, the PAM MUST respond with an AA-Answer setting the Result-Code AVP to DIAMETER_UNABLE_TO_COMPLY. Given that subscription to signaling path status is not supported, the PAM will never send a notification of signaling path status to the P-CSCF as described in [TS 29.214] clause 4.4.6.3.

7 PARAMETER MAPPING

Given that the PAM is responsible for managing the translation between the Rx and PacketCable Multimedia interfaces, this section provides requirements on how the PAM accomplishes this by defining how to map critical information between the two interfaces.

7.1 Mapping AA-Requests to Gate-Set

7.1.1 FlowSpec Traffic Profile

The Integrated Services architecture uses general purpose (layer 2 independent) descriptions of the traffic characteristics and resource requirements of a flow. The description of the traffic is known as a TSpec, the resource requirements are contained in an RSpec, and the combination of these is known as a FlowSpec. In order to reserve resources on a specific layer 2 medium such as a DOCSIS cable network, it is necessary to define a mapping from the layer two independent FlowSpec to specific layer 2 parameters. Mappings for a variety of other technologies (ATM, 802.3 LANs, etc.) have already been defined elsewhere.

Integrated Services currently defines two types of service, controlled load and guaranteed, the latter being the more suitable for latency sensitive applications. At the time of this writing, there are no requirements for the use of a controlled load service in PacketCable 2.0.

When making a reservation for guaranteed service, the FlowSpec contains:

TSpec

Bucket depth (b) - bytes

Bucket rate (r) - bytes/second

Peak rate (p) - bytes/second

Min policed unit (m) - bytes

Maximum datagram size (M) - bytes

RSpec

Reserved rate (R) - bytes/second

Slack term (S) - microseconds

The TSpec terms are mostly self-explanatory. (r,b) specifies a token bucket that the traffic conforms to, p is the peak rate at which the source will send, and M is the maximum packet size (including IP and higher layer headers) that will be generated by the source. The minimum policed unit, m, is usually the smallest packet size that the source will generate; if the source sends a smaller packet, it will count as a packet of size m for the purposes of policing.

To understand the RSpec, it is helpful to understand how delay is calculated in an Integrated Services environment. The maximum end-to-end delay experienced by a packet receiving guaranteed service is:

$$\text{Delay} = b/R + C_{\text{tot}}/R + D_{\text{tot}}$$

where b and R are as defined above, and C_{tot} and D_{tot} are accumulated 'error terms' provided by the network elements along the path, which describe their deviation from 'ideal' behavior.

The rate R provided in the RSpec is the amount of bandwidth allocated to the flow and must be greater than or equal to r from the TSpec for the above delay bound to hold. Thus, a flow's delay bound is completely determined by the choice of R ; the reason to use a value of R greater than r would be to reduce the delay experienced by the flow.

Since it is not permissible to set $R < r$, a node making a reservation may perform the above calculation and determine that the delay bound is tighter than needed. In such a case, the node may set $R = r$ and set S to a non-zero value. The value of S would be chosen such that:

$$\text{Desired delay bound} = S + b/R + C_{\text{tot}} / R + D_{\text{tot}}$$

Guaranteed Service does not attempt to bound jitter any more than is implied by the delay bound. In general, minimum delay that a packet might experience is the speed of light delay, and the maximum is the delay bound given above; the maximum jitter is the difference between these two. Thus jitter may be controlled by suitable choice of R and S .

For well known audio and video codecs, the PAM MUST use the FlowSpec parameters defined in the PacketCable CODEC and Media specification [CODEC-MEDIA].

For non-well-known codecs, the PAM MUST generate a FlowSpec via the following mechanisms:

If the $b=\text{TIAS}:\#\#$ and $a=\text{maxprate}$ parameters are supplied, the PAM MUST determine the bandwidth necessary for the session as follows:

1. Determine which lower layers will be used and calculate the sum of the sizes of the headers in bits ($h\text{-size}$). In cases of variable header sizes, the maximum size MUST be used. For RTP-transported media, the lower layers MUST include the RTP header with header extensions, if used, the CSRC list, and any profile-specific extensions;
2. Retrieve the maximum packet rate from the SDP ($\text{prate} = \text{maxprate}$);
3. Calculate the transport overhead by multiplying the header sizes by the packet rate ($t\text{-over} = h\text{-size} * \text{prate}$);
4. Round the transport overhead up to nearest integer in bits ($t\text{-over} = \text{CEIL}(t\text{-over})$);
5. Add the transport overhead to the transport independent bandwidth value (total bit-rate (B) = $\text{bw-value} + t\text{-over}$).

When the above calculation is performed using the "maxprate", the bit-rate value will be the absolute maximum the media stream may use over the transport assumed in the calculations.

If the $b=\text{TIAS}$ parameter is not supplied, the PAM MUST use the value specified in the $b=\text{AS}:\#\#$ bandwidth parameter as a substitution for the value calculated above. Please note that $b=\text{TIAS}$ is specified in bits per second, while $b=\text{AS}$ is specified in kilobits per second.

Using the calculated bandwidth parameter (B), the PAM MUST set the flowspec values to the following:

$$b = \text{Bucket Size} = (B / 8) / \text{maxprate}$$

$$p = r = R = B / 8$$

$$m = b$$

$$M = 1522 \text{ bytes}$$

7.1.1.1 Multiple Codecs

When using the SIP and SDP standards, a UE can use any of the negotiated codecs at any point during the flow. In order to inform the CMTS of the QoS requirements of the current flow and allow it to properly schedule the resources, the application must be aware of the flow parameters being used at any instant. The SIP framework, however, does not mandate this level of session awareness, i.e., SIP does not mandate that the UE notifies the signaling layer when making a change in the flow. Therefore, the PAM may not know at a certain point in time what flow parameters a UE is using in a session.

Due to the fact that the PAM does not necessarily know the flow parameters being used at any point in time on a flow, the PAM MUST commit the Least Upper Bound (LUB) when the flow is in the active state. By doing so, the PAM can be assured that the UE will always have the resources necessary for any codec they use in the negotiated set.

7.1.1.1.1 Calculation of the Least Upper Bound (LUB)

There are various situations in which a reservation needs to cover a range of possible flowspecs. For example, for some applications it is desirable to create a reservation, which can handle a switch from one codec to another mid-session without having to pass through admission control at each switch-over time.

In order to support this functionality, the PAM MUST send a flowspec which contains the Least Upper Bound (LUB) of the necessary flow parameters for the individual flows.

The Least Upper Bound (LUB) of two flows A and B, $LUB(A, B)$, is the "smallest" envelope that can carry both of the flows A, B non-simultaneously. $LUB(A, B)$ is calculated on a parameter-by-parameter basis as follows:

Define the TSpec values for a flow α . Also define the period P_α as M_α / r_α . Then $LUB(A, B)$ is given by:

$$LUB(A, B) \equiv \{b_{LUB(A, B)} \equiv \text{MAX}(b_A, b_B),$$

$$r_{LUB(A, B)} \equiv (M_{LUB(A, B)} / P_{LUB(A, B)}),$$

$$p_{LUB(A, B)} \equiv \text{MAX}(p_A, p_B, r_{LUB(A, B)}),$$

$$m_{LUB(A, B)} \equiv \text{MAX}(m_A, m_B),$$

$$M_{LUB(A, B)} \equiv \text{MAX}(M_A, M_B)$$

$$\}$$

where:

$$P_{LUB(A, B)} \equiv GCF(P_A, P_B);$$

the function $MAX(x, y)$ means "take the higher of the pair (x, y) ";

the function $MAX(x, y, z) \equiv MAX(MAX(x, y), z)$;

the function $GCF(x, y)$ means "take the Greatest Common Factor of the pair (x, y) ".

The LUB of n flows ($n \neq 2$), $LUB(n_1, n_2, \dots)$, is defined recursively as:

$$LUB(n_1, n_2, \dots, N) \equiv LUB(n_1, LUB(n_2, \dots, N))$$

In addition, the slack term in the corresponding RSpec must allow any component flow to use the resources. In order to ensure that this criterion is met, the RSpec for the flow is set to the minimum value of the RSpec values in the component flows. That is:

$$S_{LUB(A, B)} \equiv MIN(S_A, S_B)$$

Where the function $MIN(x, y)$ means take the lower of the pair (x, y) .

The following example shows how TSpec parameters are determined using LUB algorithm specified above:

1. As the result of codec negotiation, the following codecs are selected for a session:
G711(20ms) and G728(10ms)
2. The LUB bucket depth for the selected codecs is:
 $G711(20ms) = (8000 / 50) + 40 = 200$ bytes
 $G728(10ms) = (2000 / 100) + 40 = 60$ bytes
 $b[LUB] = m[LUB] = M[LUB] = MAX(200, 60) = 200$ bytes
3. The LUB bucket rate for selected codecs is:
 $P[LUB] = GCF(10ms, 20ms) = 10ms = 0.01$ second
 $r[LUB] = M * 1/P = 200 * 1/0.01 = 20,000$ bytes per second
 $r[G711(20ms)] = 200 * 1/0.02 = 10,000$ bytes per second
 $r[G728(10ms)] = 60 * 1/0.01 = 6,000$ bytes per second
 $p[LUB] = MAX(10000, 6000, 20000) = 20,000$ bytes per second

7.1.1.2 Forking

A case may exist where a session is forked and the originating UE receives several provisional responses which contain an SDP answer. Annex A of [TS 29.214] provides procedures on how the P-CSCF notifies the PCRF that forking has taken place as well as procedures on how the PCRF handles this case. While this specification does not make any changes to those procedures, it does call out some bandwidth considerations that need to be made by the PAM. In particular, if the UE has invoked the ICE procedures and advertised a media relay address (TURN server) in the m/c lines of the SDP offer, it is possible that the UE may receive traffic via the media relay (in the case of early media). Given that the UE does not yet know which provisional response will ultimately win the session request, it cannot yet set the active destination at the TURN server. Since the active destination is not set, the TURN server has to wrap each RTP packet it receives with a STUN header which identifies the source of the RTP packet prior to forwarding to the UE. This STUN header adds an additional 36 bytes of overhead plus any padding to a 4 byte boundary for the RTP data itself.

If the PAM receives an AA request within the existing Diameter session containing the SIP-Forking-Indication AVP with value SEVERAL_DIALOGUES, and include the service information derived from the latest provisional

response, forking has taken place. Further, if at least one of the AA requests related to the Diameter session has set the flow status to ENABLED-UPLINK, ENABLED-DOWNLINK, or ENABLED, media is allowed to flow between the originating UE and any of the terminating UEs. Since the UE will be unable to set the active destination at the TURN server, the PAM has to assume that each RTP packet sent to the UE contains a STUN header. In this case the PAM SHOULD include the necessary STUN header overhead when making resource allocation requests for the session (36 bytes of overhead plus any padding to a 4 byte boundary for the RTP data itself).

Once the session has been answered, all other dialogs will be cleared. Thus the PAM, upon receipt of an AA request within an existing Diameter session with no SIP-Forking-Indication AVP or with a SIP-Forking-Indication AVP with value SINGLE_DIALOGUE, the PAM SHOULD update the resource allocation to remove any previously accounted for STUN header.

7.1.1.3 Setting the Envelope Field

The FlowSpec Traffic Profile object contains an envelope field which indicates which of the Authorized, Reserved and Committed envelopes are present in the profile and therefore what state is requested for the gate. As described in [PCMM] the gate can be placed in the Authorized state by providing only the Authorized Envelope (envelope = 001), the Reserved State by providing the Authorized and Reserved envelopes (envelope = 011), or Committed state by providing all three envelopes (envelope = 111).

The PAM MUST set the envelope field based on the value of the Flow-Status AVP corresponding to the flow as follows:

- If the Flow-Status value is ENABLED-UPLINK and gate direction is downstream, then the envelope field must be set to 011 (Reserved).
- If the Flow-Status value is ENABLED-UPLINK and gate direction is upstream, then the envelope field must be set to 111 (Committed).
- If the Flow-Status value is ENABLED-DOWNLINK and gate direction is downstream, then the envelope field must be set to 111 (Committed).
- If the Flow-Status value is ENABLED- DOWNLINK and gate direction is upstream, then the envelope field must be set to 011 (Reserved).
- If the Flow-Status value is ENABLED, the envelope field MUST be set to 111 (Committed).
- If the Flow-Status value is DISABLED, the envelope field MUST be set to 011 (Reserved).

At the time of this writing, there are no requirements for the use of Authorized State in PacketCable 2.0.

7.1.2 Classifiers

The PAM MUST build PacketCable Multimedia classifiers as specified in the Flow-Description AVP.

The Flow-Description AVP defines a packet filter for an IP flow with the following information:

- Direction (in or out),
- Source and destination IP address (possibly masked),
- Protocol,
- Source and destination port (the Source Port may be omitted to indicate that any source port is allowed).

The Flow-Description AVP is used to describe a single IP flow, and the direction "in" refers to uplink IP flows, and the direction "out" refers to downlink IP flows.

The PAM MUST generate the classifier by mapping the Flow-Description AVP source and destination address source and destination port fields, and the protocol field to the appropriate PacketCable Multimedia classifier fields.

7.1.3 GateSpec

The PAM SHOULD set the DSCP/TOS Field and the DSCP/TOS mask values in the Gate Spec COPS Object based on the values received in the media type AVP of the media component description AVP of the AA Request message. The DSCP Field and mask values for each of the media types MAY be obtained from static configuration of the PAM.

The PAM SHOULD set the Session class ID value in the Gate Spec object based on the values received in the Reservation Priority and the service-URN AVPs of the AA Request message. The Session class ID values for each of the media types MAY be obtained from static configuration of the PAM.

7.1.4 SubscriberID

If the Framed-IP-Address AVP is present, the PAM MUST set the PacketCable Multimedia SubscriberID object to the same value as received in the Framed-IP-Address AVP of the AA Request message.

If the Framed-IPv6-Prefix is present, the PAM MUST set the PacketCable Multimedia IPv6SubscriberID object to the same value as received in the Framed-IPv6-Prefix AVP of the AA Request message.

7.1.5 Application Type ID

The PAM SHOULD set the two byte application type in the PacketCable Multimedia AMID object based on the value received in the AF-Application-Identifier AVP of the AA Request message. The AF-Application-Identifier AVP to Application type mapping MAY be obtained from static configuration of the PAM.

7.2 Mapping Gate-Set-Ack to AA-Answers

7.2.1 Access Network Charging Identifier

If configured to generate a BCID, the PAM MUST embed the generated BCID assigned to the resulting PacketCable Multimedia Gates in the Access-Network-Charging-Identifier AVP of the AA Answer message as a hexadecimal string. If the AA Answer is in response to an AA Request which modifies an existing DIAMETER session, the PAM MAY send the existing BCID in the AA Answer message.

Appendix I Acknowledgements

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

Chandrasekhar Annombhotla - C-Cor

Matt Tooley - CableMatrix

James Medlock - Camiant

Brian Davis - Cisco Systems

Vikrant Karan - Ellacoya Networks

David Flanagan - Motorola

Tony MacDonald - Nortel Networks

Kevin Johns - CableLabs
