

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.163

(12/2007)

SERIES J: CABLE NETWORKS AND TRANSMISSION
OF TELEVISION, SOUND PROGRAMME AND OTHER
MULTIMEDIA SIGNALS

IPCablecom

**Dynamic quality of service for the provision
of real-time services over cable television
networks using cable modems**

ITU-T Recommendation J.163

Copyright International Telecommunication Union/ITU Telecommunication Sector
Provided by IHS Markit under license with ITU-T
No reproduction or networking permitted without license from IHS

ITU-T Recommendation J.163

Dynamic quality of service for the provision of real-time services over cable television networks using cable modems

Summary

ITU-T Recommendation J.163 addresses requirements for a client device to obtain access to network resources. In particular, it specifies a comprehensive mechanism for a client device to request a specific quality of service from the DOCSIS network. Extensive examples illustrate the use of this Recommendation. The scope of this Recommendation is to define the QoS architecture for the "Access" portion of the IP-Cablecom network, provided to requesting applications on a per-flow basis. The access portion of the network is defined to be between the multimedia terminal adapter (MTA) and the cable modem termination system (CMTS), including the DOCSIS network. The method of QoS allocation over the backbone is unspecified in this Recommendation. Interface to the managed IP backbone and issues related to IP multicast are not within the scope of this Recommendation. This Recommendation also recognizes that per-flow reservations may be required within the customer premises, and the protocol developed addresses this potential need.

Source

ITU-T Recommendation J.163 was approved on 14 December 2007 by ITU-T Study Group 9 (2005-2008) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure e.g. interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2008

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Page

1	Scope	1
2	References.....	1
2.1	Normative references.....	1
2.2	Informative references.....	2
3	Terms and definitions	2
4	Abbreviations and conventions.....	2
4.1	Abbreviations	2
4.2	Conventions.....	3
5	Technical overview.....	3
5.1	IPCablecom QoS architecture requirements	5
5.2	IP QoS access network elements.....	7
5.3	IPCablecom dynamic QoS architecture.....	8
5.4	QoS interfaces	9
5.5	Framework for IPCablecom QoS	10
5.6	Requirements of access network resource management	12
5.7	Theory of operation	17
6	Embedded MTA to CM QoS Protocol (pkt-q1)	22
6.1	RSVP flowspecs	23
6.2	DOCSIS support for resource reservation.....	34
6.3	Use of DOCSIS MAC control service interface.....	40
7	Authorization interface description (pkt-q6)	44
7.1	Gates: The framework for QoS control	45
7.2	COPS profile for IPCablecom.....	50
7.3	Gate control protocol message formats	52
7.4	Gate control protocol operation.....	62
7.5	CMS use of gate protocol.....	68
7.6	Gate-coordination	68
	Annex A –Timer definitions and values	71
	Appendix I – Theft of service scenarios	73
I.1	Scenario No. 1: Customers establishing high QoS connections themselves..	73
I.2	Scenario No. 2: Customers using provisioned QoS for non-voice applications.....	74
I.3	Scenario No. 3: MTA altering the destination address in voice packets.....	74
I.4	Scenario No. 4: Use of half-connections	74
I.5	Scenario No. 6: Early termination leaving a half-connection	74
I.6	Scenario No. 6: Forged gate coordination messages.....	74
I.7	Scenario No. 7: Fraud directed against unwanted callers.....	75

	Page
Appendix II – COPS (common open policy service)	76
II.1 COPS procedures and principles	76
II.2 Comparison of COPS and LDAP for policy	77
Appendix III – TCP considerations	78
III.1 Requirements	78
III.2 Recommended changes	78
III.3 TCP connection establishment impacting post-dial delay.....	79
III.4 Need for low latency for packets between the GC and CMTS, even under loss.....	79
III.5 Head-of-line blocking.....	80
III.6 TCP slow start	80
III.7 Delaying of packets: Nagle's algorithm.....	81
III.8 Non-blocking interface	81

ITU-T Recommendation J.163

Dynamic quality of service for the provision of real-time services over cable television networks using cable modems

1 Scope

This Recommendation addresses requirements for a client device to obtain access to network resources. In particular, it specifies a comprehensive mechanism for a client device to request a specific quality of service from the DOCSIS network. Extensive examples illustrate the use of this Recommendation. The scope of this Recommendation is to define the QoS architecture for the "Access" portion of the IP-Cablecom network, provided to requesting applications on a per-flow basis. The access portion of the network is defined to be between the multimedia terminal adapter (MTA) and the cable modem termination system (CMTS), including the DOCSIS network. The method of QoS allocation over the backbone is unspecified in this Recommendation. Interface to the managed IP backbone and issues related to IP multicast are not within the scope of this Recommendation. This Recommendation also recognizes that per-flow reservations may be required within the customer premises, and the protocol developed addresses this potential need.

NOTE – The structure and content of this Recommendation have been organized for ease of use by those familiar with the original source material; as such, the usual style of ITU-T recommendations has not been applied.

2 References

2.1 Normative references

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- [ITU-T J.83] ITU-T Recommendation J.83 (2007), *Digital multi-programme systems for television, sound and data services for cable distribution*.
- [ITU-T J.112] ITU-T Recommendation J.112 (1998), *Transmission systems for interactive cable television services*.
- [ITU-T J.112-A] ITU-T Recommendation J.112 Annex A (2001), *Digital Video Broadcasting: DVB interaction channel for Cable TV (CATV) distribution systems*.
- [ITU-T J.112-B] ITU-T Recommendation J.112 Annex B (2004), *Data-over-cable service interface specifications: Radio-frequency interface specification*.
- [ITU-T J.160] ITU-T Recommendation J.160 (2005), *Architectural framework for the delivery of time-critical services over cable television networks using cable modems*.
- [ITU-T J.161] ITU-T Recommendation J.161 (2007), *Audio and video codec requirements and usage for the provision of bidirectional audio services over cable television networks using cable modems*.
- [IETF RFC 2748] IETF RFC 2748 (2000), *The COPS (Common Open Policy Service) Protocol*.

2.2 Informative references

- [ITU-T G.114] ITU-T Recommendation G.114 (2003), *One-way transmission time*.
- [ITU-T G.711] ITU-T Recommendation G.711 (1988), *Pulse code modulation (PCM) of voice frequencies*.
- [ITU-T G.726] ITU-T Recommendation G.726 (1990), *40, 32, 24, 16 kbit/s adaptive differential pulse code modulation (ADPCM)*.
- [ITU-T G.728] ITU-T Recommendation G.728 (1992), *Coding of speech at 16 kbit/s using low-delay code excited linear prediction*.
- [ITU-T G.729] ITU-T Recommendation G.729 Annex E (1998), *11.8 kbit/s CS-ACELP speech coding algorithm*.
- [ITU-T J.162] ITU-T Recommendation J.162 (2007), *Network call signalling protocol for the delivery of time-critical services over cable television networks using cable modems*.
- [ITU-T J.164] ITU-T Recommendation J.164 (2007), *Event message requirements for the support of real-time services over cable television networks using cable modems*.
- [ITU-T J.170] ITU-T Recommendation J.170 (2005), *IPCablecom security specification*.
- [IETF RFC 791] IETF RFC 791 (1981), *Internet Protocol – DARPA Internet Program – Protocol specification*.
- [IETF RFC 2327] IETF RFC 2327 (1998), *SDP: Session Description Protocol*.
- [IETF RFC 2474] IETF RFC 2474 (1998), *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*.
- [IETF RFC 2753] IETF RFC 2753 (2000), *A Framework for Policy-based Admission Control*.
- [IETF RFC 3551] IETF RFC 3551 (2003), *RTP Profile for Audio and Video Conferences with Minimal control*.

3 Terms and definitions

This Recommendation defines the following terms:

3.1 cable modem: A cable modem is a layer two termination device that terminates the customer end of the J.112 (or J.122) connection.

3.2 DOCSIS flow: A unidirectional or bidirectional flow of data packets that is subject to MAC-layer signalling and QoS assignment compliant to [ITU-T J.112] (or ITU-T Rec. J.122).

3.3 IPCablecom: An ITU-T project that includes an architecture and a series of Recommendations that enable the delivery of real-time services over the cable television networks using cable modems.

4 Abbreviations and conventions

4.1 Abbreviations

This Recommendation uses the following abbreviations:

CM Cable Modem

CMTS Cable Modem Termination System

COPS Common Open Policy Service

CPE	Customer Premises Equipment
DCS	Distributed Call Signalling
DSA	Dynamic Service Addition
DSC	Dynamic Service Change
INA	Interactive Network Adapter
IP	Internet Protocol
MTA	Multimedia Terminal Adapter
NCS	Network-based Call Signalling
PHS	Payload Header Suppression
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RAP	Resource Allocation Protocol
RSVP	Resource ReSerVation Protocol
TLV	Type-Length-Value
VAD	Voice Activity Detection

4.2 Conventions

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

"MUST"	This word or the adjective "REQUIRED" means that the item is an absolute requirement of this Recommendation.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this Recommendation.
"SHOULD"	This word or the adjective "RECOMMENDED" 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 behaviour is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behaviour described with this label.
"MAY"	This word or the adjective "OPTIONAL" 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.

5 Technical overview

Enhanced quality of service is required for supporting interactive multimedia applications. Resources may be constrained in segments of the network, requiring allocation of resources in the network. The scope of this Recommendation is to define the Quality of Service architecture for the "Access" portion of the IP-Cablecom network. The access portion of the network is defined to be between the multimedia terminal adapter (MTA) and the cable modem termination system (CMTS), including the DOCSIS network. This Recommendation also recognizes that per-flow reservations

may be required within the customer premises, and the protocols developed herein address this potential need. Although some segments of the backbone network may require resource reservation to provide adequate quality of service, we consider the protocols for backbone resource management to be outside the scope of this Recommendation.

Resources are allocated on the DOCSIS network for individual flows associated with each session of an application, per subscriber, on an authorized and authenticated basis. A DQoS session, or simply a session, is defined by this Recommendation to be a single bidirectional data flow between two clients. When a multimedia application needs multiple bidirectional data flows (e.g., one for voice and a separate for video), separate DQoS sessions are established for each. Applications may use only half of the session's bidirectional data flow, thereby providing send-only or receive-only services. For example, in a typical voice communications application, a simple communication between two parties is implemented by a single session, while complex, multiparty communications (e.g., "conference calls") are implemented by multiple simultaneous sessions.

Network-based call signalling [ITU-T J.162] is the defined IP-Cablecom call signalling protocol. This Dynamic QoS specification is the underlying QoS framework for this call signalling protocol. QoS is allocated for flows associated with a session in concert with the signalling protocol.

This Recommendation introduces the concept of a segment-by-segment QoS framework. It exploits the information available from signalling protocols to perform the QoS assignment on both the "local" segment (on the DOCSIS network close to the originating party) and the "remote" segment (the DOCSIS network close to the terminating party). Thus, this Recommendation allows different providers to use the most appropriate mechanisms for the segment that they are managing. Using a concatenation of the segments with QoS, we provide end-to-end QoS assurance for the session.

The Dynamic QoS specification incorporates protocols to enable providers of packet-based voice communications using the IP-Cablecom framework to use different charging models, including both flat-rate charging as well as usage-based charging. It is the intent of this Recommendation to ensure that enhanced QoS is provided only to authorized and authenticated users. The specific techniques used for authorizing and authenticating a user are beyond the scope of this Recommendation.

This Dynamic QoS specification recognizes the requirements of a commercially viable voice communications service analogous to that offered by means of the public switched telephone network. It is important to ensure that resources are available before the two parties involved in the session are invited to communicate. Thus, resources are reserved before the recipient of the communication is notified that someone is trying to initiate a communication. If there are insufficient resources for a session, then the session is blocked.

The protocols developed in this Recommendation explicitly recognize the need to ensure that there is no potential for fraud or theft of service by endpoints that do not wish to cooperate with the call signalling and QoS signalling protocols with the intent of avoiding being charged for usage. This Recommendation introduces the concept of a two-phase for resource reservation (reserve and commit). The two phases allow a provider to both allocate resources only when they are required (when the voice path is cut-through) which may be used for billing. Further, because the second phase to commit resources requires an explicit request from the MTA, it enables the provider to prevent fraud and theft of service.

This Recommendation is technically compatible with the corresponding CableLabs PacketCable document: *PacketCable Dynamic Quality-of-Service Specification* PKT-SP-DQOS1.5 I01.

5.1 IPCablecom QoS architecture requirements

The following list presents the QoS requirements for supporting multimedia applications over IPCablecom Networks.

1) *Provide IPCablecom accounting for the QoS resources on a per-session basis*

It is anticipated that, from a billing perspective, one of the resources that will need to be accounted for is the use of QoS in the DOCSIS network. Thus, information needs to be identified and tracked that allows reconciliation of the use of the DOCSIS QoS resource with IPCablecom session activity.

2) *Both two-phase (reserve-commit) and single-phase (commit) QoS activation models*

Under application control, it should be possible to utilize either a two-phase or single-phase QoS activation model. In the two-phase model the application reserves the resource, then later commits it. In the single-phase model, both reservation and commitment occur as a single autonomous operation. As in the DOCSIS model, resources that are reserved but not yet committed are available for temporary assignment to other (e.g., best effort) service flows. This Recommendation provides mechanisms for both two-phase and single-phase activation for embedded MTAs.

3) *Provide IPCablecom defined policies to control QoS in both the DOCSIS network and the IP backbone*

It should be possible for different types of sessions to have different QoS characteristics. For example, sessions within a single Cable Operator provider's domain may receive different QoS than sessions outside the domain (e.g., international sessions including links to the PSTN). This dynamic QoS specification may allow a Cable Operator to provide different QoS for different types of customers (e.g., higher QoS for subscribers of a business service at certain times of the day compared to residential customers), or different types of applications for a single customer.

4) *Prevent (minimize) abusive QoS usage*

Two types of abusive QoS usage are identified: that which is accurately billed but leads to denying service to others, and that which is not accurately billed and leads to theft of service. Subscriber applications and IPCablecom applications (either embedded or PC-based) may inadvertently or intentionally abuse their QoS privileges (e.g., use of enhanced QoS, which the provider wants limited to voice applications, by an FTP application). Even though the DOCSIS network is expected to enforce a subscriber's access to QoS, rich packet classification and signalling control mechanisms should exist to keep the subscriber (and the subscriber devices) from fraudulent use of QoS. Admission control procedures should be employed to reduce denial-of-service attacks.

5) *Provide admission control mechanisms for both upstream and downstream directions in the DOCSIS network*

Both upstream and downstream QoS should be subject to per-session admission control.

6) *DOCSIS QoS*

It should be possible to police (defined as marking, dropping, or delaying packets) all aspects of QoS defined in the service at the CMTS using the DOCSIS QoS mechanisms. Furthermore, it should be possible to support multiple flow mapping models: associate a single IPCablecom session to a single Service Flow and multiple IPCablecom sessions to a single Service Flow.

- 7) *Policy is enforced by the CMTS*
- Ultimate policy control is entrusted to the CMTS. The philosophy is that any client can make any QoS request, but the CMTS (or an entity behind the CMTS) is the only entity entrusted to grant or deny QoS requests.
- 8) *IPCablecom entities must be as unaware as possible of specific DOCSIS QoS primitives and parameters*
- For IPCablecom, like any other application that uses the IP-network, the design objective is to minimize the amount of access-link-specific knowledge contained within the application layer. The less access-link knowledge in the application layer, the more applications will be available for development and deployment, and the fewer testing and support-problems will be encountered.
- 9) *Reclamation of QoS resources for dead/stale sessions*
- It is necessary to reclaim and reallocate precious QoS resources for sessions that are no longer active, but have not been properly torn down. There should be no resource "leaks" in the DOCSIS link. For example, if an IPCablecom client module malfunctions in the midst of an IPCablecom session, all DOCSIS QoS resources used by the session should be released within a reasonable period of time.
- 10) *Dynamic QoS policy changes*
- It is desirable to dynamically change QoS policies for subscribers. For example, this requirement addresses the ability to change a customer's service level (e.g., upgraded from a "bronze" service to a "gold" service) on-the-fly without resetting the CM.
- 11) *Absolute minimum session set-up latency time and post pick-up delay*
- The IPCablecom Network should allow for emulation and enhancement of the PSTN experience to the user, and should be equally good, if not better, in session set-up and post pick-up delay metrics.
- 12) *Multiple concurrent sessions*
- It is desirable to allocate QoS resources (e.g., bandwidth) for not only individual point-to-point sessions, but also for multiple point-to-point sessions (e.g., conference calls, combined audio/video calls).
- 13) *Dynamic adjustment of QoS parameters in the middle of IPCablecom sessions*
- It should be possible for the IPCablecom service to change QoS mid-session, e.g., network-wide resource adjustments or creation of compatible CODEC parameters (necessitating QoS changes), or user defined feature to vary QoS levels, or detection of fax or modem streams (necessitating change from compressing CODEC to [ITU-T G.711]).
- 14) *Support multiple QoS control models*
- Strong cases can be made for both subscriber-side and network-side initiation of QoS signalling. In subscriber side signalling, an application can initiate its request for QoS immediately when the application believes it needs QoS. Also, subscriber side signalling supports application models that are peer-to-peer. In network-side signalling, implementation of the endpoint application can be completely unaware of QoS (especially in the DOCSIS network). Network-side signalling supports application models that are client-server (with the server being trusted). It is expected that both models will be present in IPCablecom (and other application) networks. This Recommendation is for subscriber-side signalling only.

15) *Support both embedded-MTA and stand-alone-MTA QoS signalling*

It should be possible to signal QoS from both an embedded-MTA and stand-alone-MTA. This Recommendation only covers the embedded MTA using direct access to the DOCSIS MAC signalling.

5.2 IP QoS access network elements

The following network elements are employed to support QoS for IPCablecom Networks.

5.2.1 Multimedia Terminal Adapter (MTA)

The IPCablecom network client device (i.e., the MTA) can be one of the following devices. These devices reside at the customer site and are connected through the DOCSIS channel to the network. All MTAs are assumed to implement some multimedia signalling protocol, such as [ITU-T J.162]. An MTA may be either a device with a standard two-wire telephone set in the MTA-1 configuration, or may add video input/output capabilities in the MTA-2 configuration. It may have minimal capabilities, or may implement this functionality on a multimedia personal computer, and have all of the capabilities of the PC at its disposal.

From the point of view of QoS, there are two types of MTAs.

- 1) **Embedded/Integrated MTA:** This is a client multimedia terminal which incorporates a DOCSIS MAC-layer interface to the DOCSIS network.
- 2) **Stand-alone MTA:** This is a Client that implements the multimedia functionality without incorporating a DOCSIS MAC-layer interface. The stand-alone MTA will typically use Ethernet, USB, or IEEE 1394 as the physical interconnect to a CM. The stand-alone MTA may be connected to a customer network, and use transport facilities of the customer network (possibly including intermediate IP routers) to establish sessions over the DOCSIS network.

5.2.2 Cable modem (CM)

This is an IPCablecom network element as defined by [ITU-T J.112], ITU-T Rec. J.122 or later Recommendations. The CM is responsible for classifying, policing and marking packets once the traffic flows are established by the signalling protocols described herein.

5.2.3 Cable modem termination system (CMTS)

The CMTS is responsible for allocating and scheduling upstream and downstream bandwidth in accordance with MTA requests and QoS authorizations established by the network administrator. The CMTS acts as a policy enforcement point (PEP) per the IETF resource allocation protocol (RAP) framework [IETF RFC 2753].

The CMTS implements an "IPCablecom Dynamic QoS Gate" (hereafter called just "Gate") between the DOCSIS network and an IP Backbone. The Gate is implemented using the packet classification and filtering functions defined in [ITU-T J.112] and ITU-T Rec. J.122.

The CMTS may or may not also be configured as an "IS-DS Boundary" entity. An IS-DS Boundary interfaces to an internetwork using the integrated services (IntServ) model of QoS control and some other model, e.g., differentiated services (DiffServ).

5.2.4 Call management server (CMS) and gate controller (GC)

The IPCablecom call management server (CMS) entity performs services that permit MTAs to establish multimedia sessions (including voice communications applications such as "IP telephony" or "VoIP"). The term gate controller (GC) is used to refer to the portion of either type of CMS that performs the quality of service related functions.

In the IPCablecom dynamic QoS model, the gate controller controls the operation of the Gates implemented on a CMTS. The GC acts as a policy decision point (PDP) per the IETF resource allocation protocol (RAP) framework [IETF RFC 2753].

5.2.5 Record keeping server (RKS)

The record keeping server is an IPCablecom network element that only receives information from IPCablecom elements described in this Recommendation. The RKS can be used as a billing server, diagnostic tool, etc.

5.3 IPCablecom dynamic QoS architecture

The IPCablecom QoS architecture is based upon [ITU-T J.112], IETF RSVP, and IETF integrated services guaranteed QoS.

Specifically, the IPCablecom QoS architecture uses the protocol as defined in [ITU-T J.112] within the cable television network. These messages support static and dynamic installation of packet classifiers (i.e., Filter-Specs) and flow scheduling (i.e., flowSpecs) mechanisms to deliver enhanced quality of service. DOCSIS QoS is based upon the objects which describe traffic and flow specifications, similar to the TSpec and RSpec objects as defined in the IETF resource reservation protocol (RSVP). This allows QoS resource reservations to be defined on a per-flow basis.

In the DOCSIS QoS architecture, traffic flows are considered as unidirectional – thus an interactive session comprises two flows, each subject to the operations shown below. For each (unidirectional) flow:

The CM, where traffic enters the IP QoS enabled cable network, is responsible for:

- Classification of IP traffic into IP QoS flows based on defined filter specifications.
- Performing traffic shaping and policing as required by the flow specification.
- Maintaining state for active flows.
- Altering the TOS field in the upstream IP headers based on the network operator's policy.
- Obtaining the required QoS from the CMTS.
- Applying DOCSIS QoS mechanisms appropriately.

The CMTS is responsible for:

- Providing the required QoS to the CM based upon policy configuration.
- Allocating upstream bandwidth in accordance with CM requests and network QoS policies.
- Classifying each arriving packet from the network side interface and assigning it to a QoS level based on defined filter specifications.
- Policing the TOS field in received packets from the cable network to enforce TOS field settings per network operator policy.
- Altering the TOS field in the downstream IP headers based on the network operator's policy.
- Performing traffic shaping and policing as required by the flow specification.
- Forwarding downstream packets to the DOCSIS network using the assigned QoS.
- Forwarding upstream packets to the backbone network devices using the assigned QoS.
- Maintaining state for active flows.

The backbone network may either utilize IETF integrated services-based mechanisms or use IETF differentiated services mechanisms. In a DiffServ backbone, network routers forward a packet, providing the appropriate IP QoS, based on the setting of the TOS field. In a DiffServ backbone, no per-flow state is required in the core network devices.

5.4 QoS interfaces

Quality of service signalling interfaces are defined between many of the components of the IPCablecom network as shown in Figure 1. Signalling involves communication of QoS requirements at the application layer (e.g., SDP parameters), network layer (e.g., RSVP), and at the data-link layer (e.g., DOCSIS QoS). Also, the requirement for policy enforcement and system linkages between the OSS subscriber provisioning, admission control within the managed IP backbone, and admission control within the DOCSIS network creates the need for additional interfaces between components in the IPCablecom network.

An expanded explanation of QoS architecture framework is contained in the IPCablecom architecture framework, [ITU-T J.160], and is shown in Figure 1.




Figure 1 – QoS signalling interfaces in IPCablecom network

Interfaces pkt-q1 through pkt-q7 are available for controlling and processing QoS. Not all interfaces are used in all configurations and protocol variations. All but the pkt-q5 interface are utilized by DQoS. Table 1 briefly identifies each interface and how each interface is used in this dynamic QoS specification (DQoS).

Table 1 – DQoS interfaces

Interface	Description	DQoS embedded MTA (optional)
pkt-q1	MTA-CM	E-MTA MAC-Control Service Interface
pkt-q2	CM-CMTS	DOCSIS QoS, CM-initiated
pkt-q3	MTA-GC/CMS	NCS
pkt-q4	CM-Provisioning Server	N/A
pkt-q5	GC-CMTS	Gate Management
pkt-q6	CMTS-RKS	Billing
pkt-q7	CMS-CMS	CMS-to-CMS Signalling

pkt-q1: Interface between the MTA and CM

This interface is only defined for the embedded MTA. The interface decomposes into three sub-interfaces:

- Control: used to manage DOCSIS service flows and their associated QoS traffic parameters and classification rules.
- Synchronization: used to synchronize packetization and scheduling for minimizing latency and jitter.
- Transport: used to process packets in the media stream and perform appropriate per-packet QoS processing.

This interface is conceptually defined in [ITU-T J.112]. For stand-alone MTAs, no instance of this interface is defined.

pkt-q2: DOCSIS QoS interface between CM and CMTS

This is the DOCSIS QoS interface (control, scheduling and transport). Control functions can be initiated from either the CM or the CMTS. However, the CMTS is the final policy arbiter and granter of resources by performing admission control for the DOCSIS network. This interface is defined in [ITU-T J.112].

pkt-q3: Application layer signalling between GC/CMS and MTA

This is the network-based call signalling protocol [ITU-T J.162]. Many parameters are signalled across this interface such as the media stream, IP addresses, port numbers, and the selection of Codec and packetization characteristics. This interface is defined in [ITU-T J.162].

pkt-q4: Signalling from the DOCSIS/IPCablecom provisioning to the CM

This interface is not utilized for QoS signalling in DQoS.

pkt-q5: Interface between the GC/CMS and CMTS

This interface is used to manage the dynamic Gates for media stream sessions. This interface enables the IPCablecom network to request and authorize QoS.

pkt-q6: CMTS to record keeping server

This interface is used by the CMTS to signal to the RKS all changes in session authorization and usage.

pkt-q7: CMS to CMS interface

This interface is used for session management and resource coordination between a pair of CMSs.

5.5 Framework for IPCablecom QoS

In order to justify its costs to the end user, a commercial multimedia service (e.g., voice communications capability) may require a high level of transport and signalling performance, including:

- Low delay: End-to-end packet delay needs to be small enough that it does not interfere with normal multimedia interactions. For normal telephony service using the PSTN, the ITU-T recommends no greater than 300 ms round-trip delay¹. Given that the end-to-end backbone

¹ [ITU-T G.114] states that a one-way delay of 150 ms is acceptable for most user applications. However, highly interactive voice and data applications may experience degradation even when delays are below 150 ms. Therefore, any increase in processing delay (even on connections with transmission times well below 150 ms) should be discouraged unless there are clear service and application benefits.

propagation delay may absorb a significant amount of this delay budget, it is important to control delay on the access channel, at least for long-distance calls.

- Low packet loss: Packet loss needs to be small enough so that voice quality or performance of fax and voiceband modems is not perceptibly impaired. While loss concealment algorithms can be used to reproduce intelligible speech even with high loss rates, the resulting performance cannot be considered to be adequate as a replacement for existing circuit-switched telephone service. Loss requirements for acceptable voiceband modem performance are even more stringent than those for voice.
- Short post-dial delay: The delay between the user signalling a connection request and receiving positive confirmation from the network needs to be short enough that users do not perceive a difference from the post-dial delay they are accustomed to in the circuit switched network, or believe that the network has failed. This is of the order of one second.
- Short post pick-up delay: The delay between a user picking up a ringing phone and the voice path being cut through needs to be short enough so that the "hello" is not clipped. This should be less than a few hundred milliseconds (ideally less than 100 ms).

A key contribution of the Dynamic QoS framework is a recognition of the need for coordination between signalling, which controls access to application specific services, and resource management, which controls access to network-layer resources. This coordination provides a number of critical functions. It ensures that users are authenticated and authorized before receiving access to the enhanced QoS associated with the service. It ensures that network resources are available end-to-end before alerting the destination MTA. Finally, it ensures that the use of resources is properly accounted for, consistent with the conventions of traditional voice-grade telephone service (to which some IP-Cablecom services are similar from a customer perspective) in which charging occurs only after the party receiving a communication picks up.

In order to support the above requirements, the QoS protocols assure that all resources are committed to all transport segments before the signalling protocols cause alerting of the destination. Likewise, during tear-down of a session, the QoS protocols include measures to assure that all resources dedicated exclusively to the session are released. Without this coordination between the two directions of data flows, it would be possible for users to thwart the QoS controls and obtain free service. For example, if the paying client terminates the session, but the non-paying does not, a "half channel" remains that can be used to fraudulently transfer data in one direction. The QoS protocols approximate the "all or nothing" transaction semantics for session creation and destruction.

It is desired that the mechanisms used to implement the session be based on existing standards and practices, and also that the results of this work be usable to support alternative call models. These desires have led to the use of the IETF real time protocol (RTP) to carry multimedia data, carried over the IETF user datagram protocol (UDP). In-band signalling to set up quality of service is carried out using DOCSIS Dynamic QoS messages.

The QoS architecture should provide support for new emerging applications that are dependent on multicast data delivery. Although this is not a strict requirement in the QoS architecture, providing support for multicast will enable the future development of a rich set of multimedia applications. It has not yet been examined whether the resource management enhancements introduced here will support multicast seamlessly or not.

For purposes of managing quality of service, the bearer channel for a session is managed as three distinct segments: the access network for the originating side of the session, a backbone network, and the access network for the terminating side of the session. DOCSIS network resources are managed as a pair of dynamic service flows using the mechanisms defined in [ITU-T J.112]. Backbone resources may be managed either per-flow or, more likely, through an aggregated quality

of service mechanism. Management of backbone resources is outside the scope of this Recommendation.

A QoS-defined construct called a *gate* provides a control point for the connection of access networks to high quality backbone service. A gate is implemented by a CMTS and consists of a packet classifier, a traffic policer, and an interface to an entity that gathers statistics and events (all of these components exist in the DOCSIS network). A gate can ensure that only those sessions that have been authorized by the service provider receive high quality service. Gates are managed selectively for a flow. For IP-Cablecom-based voice communications service, they are opened for individual calls. Opening a gate involves an admission control check that is performed when a resource management request is received from the client for an individual session, and it may involve resource reservation in the network for the session if necessary. The upstream packet filter in the gate allows a flow of packets to receive enhanced QoS for a session from a specific IP source address and port number to a specific IP destination address and port number. The downstream packet filter in the gate allows a flow of packets to receive enhanced QoS for a session from a specific IP source address to a specific IP destination address and port number.

A Gate is a logical entity that resides in a CMTS. A GateID is associated with an individual session and is meaningful at the Gate; the GateID is an identifier that is locally unique at the CMTS, and is assigned by that CMTS. A Gate is unidirectional in nature. If a Gate is "Closed", then data going upstream/downstream on the DOCSIS access network may either be dropped or provided best-effort service. The choice of dropping packets or serving them on a best-effort basis is a policy choice of the provider.

The gate controller is responsible for the policy decision of when and whether the gate should be opened. A gate is established in advance of a resource management request. This allows the policy function, which is at the gate controller, to be "stateless" in that it does not need to know the state of sessions that are already in progress.

While the gate controls the QoS-guaranteed stream, other flows, such as RTCP or signalling messages, are not policed by the gate. The support of enhanced QoS for signalling messages may play a very important role if the cable system is utilizing high best-effort data traffic. In order to meet the signalling performance targets given at the beginning of this clause, it may be crucial to use a dedicated signalling flow with proper QoS constructs. It should further be noted that the exact nature of the QoS that should be given to the dedicated signal flow depends on traffic and the CMTS design and is left as a vendor differentiation point.

5.6 Requirements of access network resource management

Providing voice communications service over IP networks with the same level of quality as is available over the PSTN imposes bounds on loss and delay metrics for voice packets and requires active resource management in both the access and backbone networks. The service provider needs to be able to control access to network resources, in order to ensure that adequate capacity is available on an end-to-end basis, even under unusual or overload conditions. The service provider may seek additional revenue for providing a voice communications service with these enhanced quality characteristics (i.e., quality beyond that obtained with a "best-effort" service). The mechanisms provided herein for managed access to enhanced QoS enable the service provider to ensure that access is provided only to authorized and authenticated users on a session-by-session basis and there is no theft of that service.

Clients of the service signal their traffic and performance parameters to the "gate" at the network edge, where the network makes an admission control decision based on both resource availability as well as policy information associated with the gate.

In DOCSIS networks capacity is limited and it is necessary to do resource management on a per-flow basis. In the backbone there may be several alternatives, ranging from per-flow per-hop

admission control to coarse-grained resource provisioning. This Recommendation deals only with access network QoS, and is agnostic about backbone network QoS schemes.

5.6.1 Preventing theft of service

The network resources dedicated to the session are protected from misuse, including:

- **Authorization and Security:** Ensuring that users are authenticated and authorized before receiving access to the enhanced QoS associated with the voice communications service. The CMS/Gate Controller involved in call signalling is trusted to perform these checks and is the only entity which is trusted to create a new gate in a CMTS. The CMS/GC acts as a policy decision point from the perspective of QoS management.
- **Resource control:** Ensuring that the use of resources is properly accounted for, consistent with the conventions of providers that are part of the PSTN in which charging occurs only after the called party picks up. This includes prevention of utilization of reserved resources for purposes other than the session to which they are assigned. This is achieved through the use of gates and coordination between gates, which bind together address filtering mechanisms with resource reservations.

Since this service may be billed on a per-use basis, there is a significant risk of fraud and service theft. The architecture enables the provider to charge for quality of service. Thus, it prevents theft of service scenarios, several of which are described in Appendix I.

Theft of service scenarios are addressed in this Recommendation and other Recommendations. They motivate some of the components of the QoS and call signalling architectures and protocols.

5.6.2 Two-phase resource commitment

A two-phase protocol for resource commitment is essential to a commercial-grade voice communications service, for two reasons unique to the requirements associated with such a service. First, it ensures that resources are available before signalling the party at the far end that a communication is incoming. Secondly, it ensures that usage recording and billing are not started until the far end picks up, which is also the point at which voice may be cut-through. These properties are provided by conventional telephony signalling protocols; the same semantics will be emulated here. Also, if bandwidth is allocated before the far end picks up, a theft of service becomes possible. Requiring the endpoints to explicitly send a commitment message ensures that usage recording is based on knowledge of the endpoint and its explicit action.

This framework also supports entities, such as announcement servers and PSTN gateways, that need the voice to be cut through after the first phase of the resource management protocol.

5.6.3 Segmented resource assignment

The dynamic QoS architecture partitions resource management into distinct access and backbone segments. Segmented resource assignment is beneficial for two reasons:

- It allows for different bandwidth provisioning and signalling mechanisms for originator's network, far-end network, and backbone network.
- It allows for resource-poor segments to maintain per-flow reservations and carefully manage resource usage. At the same time, when backbone segments have sufficient resources to manage resources more coarsely, it allows the backbone to avoid keeping per-flow state, and thus enhance scalability.

When the backbone does not require explicit per-flow signalling (such as with a Diffserv backbone), it reduces the time taken to set up a session (minimize post-dial delay) and avoids impacting the voice cut-through time (minimize post-pick-up delay).

It potentially reduces the amount of reservation state that is stored if the remote client is a PSTN gateway.

After the first phase of call signalling, both clients have completed capability negotiation and know what resources are needed end-to-end. Clients send resource management messages the MAC control services interface. The CMTS maps the resource management messages to the resource management protocol used over the backbone (e.g., IETF DiffServ). It also maps the resource management message to the resource management protocol used over the access link (i.e., DOCSIS).

5.6.4 Resource changes during a session

It is possible to change the resources allocated for a session during the life of the session. This facilitates mid-session changes such as switching from a low-rate voice codec to G.711 when modem tones are detected, and the addition of video data to a session that starts as voice only.

5.6.5 Dynamic binding of resources

Dynamic binding of resources (re-reserve) is a requirement to enable efficient use of resources when services such as call waiting are invoked. Abstractly, re-reserving takes bandwidth allocated for a session between a VoIP host and a client and reallocates that same bandwidth to a session with a different client.

It is important to understand the potential danger in de-allocating the session bandwidth, then making a new request for allocation of the new bandwidth. There is a risk of another client using the last remaining bandwidth between the two steps, leaving the original session without an assured quality path. The one-step re-reserve mechanism avoids this, as the bandwidth is not made available to other clients.

5.6.6 Dynamic QoS performance

QoS messaging takes place in real time while callers wait for services to be activated or changed. Thus, the protocol needs to be fast. The number of messages is minimized, especially the number of messages which transits the backbone, and the number of upstream DOCSIS messages.

DOCSIS management messages, and call signalling messages (collectively referred to as signalling messages) are all transported over the DOCSIS network on a best-effort basis. If the CM is also supporting data services, best-effort service may be unable to provide the low latency needed for signalling messages. In this situation, the CM MAY be provisioned with a separate service flow, with enhanced QoS, to carry signalling traffic. For example, the signalling service flow could use real-time polling, or non-real-time polling service. This separate service flow is provisioned in the same manner as other DOCSIS media streams, and MAY include classifiers such that its presence is transparent to the MTA.

5.6.7 Session class

Resources may be reserved for different types of services and each service may in turn define different classes of services for its sessions. QoS reservations for sessions designated by the service provider to be of higher priority (e.g., emergency calls) suffer a lower likelihood of blocking than normal sessions. The determination of what session class to assign to a session is performed by the service provider, and is a policy that is exercised by the originating call agent/gate controller complex at the time the initial session request.

5.6.8 Intermediate network support

The architecture should not prohibit intermediate networks between the MTA or Multimedia host and the CM (e.g., customer network). Although the intermediate network may not fall under the Cable Operator's administrative domain or responsibility, allocation of bandwidth in the Cable Operator's DOCSIS network is possible when an intermediate network exists. It is also desirable to present a solution that transparently allows for the reservation of resources on the intermediate network.

5.6.9 Backbone QoS support

It is possible that some mechanism for explicitly managing backbone resources will be necessary. The scope of this Recommendation is QoS over the DOCSIS network, but the architecture provides open, sufficiently general interfaces that are compatible with many of the known backbone QoS mechanisms.

5.6.10 Handling multiple codecs

The NCS signalling used within IPCablecom allows for connections to be established with multiple codecs. In the case where a connection is successfully negotiated with multiple codecs in the list, it is important that the proper resources are allocated to make subsequent codec changes within the negotiated list work as expected. However, it is up to the CMS as to when it authorizes bandwidth during the call set-up phase, it is also under the CMS control as to how efficient it would like to be in its Authorized envelop. Should it choose to authorize bandwidth before the initial NCS CreateConnection command (CRCX), it would need to base the Authorized envelop on the proposed LCO parameters (since it does not know the subset that the MTA might negotiate). If the CMS waits until later in the call set-up phase when the codecs have been further negotiated, then it could authorize a subset of the LCO based on the current negotiated list without any negative impact (the DSA/DSC will still pass authorization). Here are the resource components that need to be allocated:

- Authorized bandwidth: When the CMS asks the MTA to reserve or commit resources by including a GateID in an NCS CreateConnection or ModifyConnection command (CRCX or MDCX), the CMS MUST ensure that the Authorized bandwidth in the gate will handle any legal resource request (DSA/DSC) from the MTA to the CMTS that results from the codec negotiation procedure. In other words, the bandwidth authorized by the CMS/GC MUST be greater than or equal to the Least-Upper-Bound of the negotiated codec list.
- Reserved bandwidth: The MTA MUST reserve the least-upper-bound of the codec bandwidth that can be used during the call (possible codecs are determined from codec negotiation procedure defined in clause 6.7 of [ITU-T J.162]).
NOTE – If the Reserved bandwidth is greater than the committed bandwidth, then the reserved bandwidth needs to be refreshed using DSCs to the CMTS.
- Committed bandwidth: The MTA MUST only commit the current codec in use in the upstream direction. This allows the extra unused bandwidth (difference between the Reserved and Committed) to be used for best-effort traffic. In the downstream direction, the MTA MUST commit the Least-Upper-Bound of the codec bandwidth that can be used during the call (possible codecs are determined from codec negotiation procedure defined in clause 6.7 of [ITU-T J.162]).

This procedure ensures that a CMS request to switch to any one of the codecs in the negotiated list will be successful. This is especially important in supporting features such as fax/modem that require a switch to G.711 for successful transmission.

If a system provider feels that the above allocation of resources places too much of a constraint on the number of voice channels that can be supported (since resources may be over reserved in many cases), then the CMS only needs to state a single codec in the LocalConnectionOptions of the connection request. This will ensure that the reserved and committed resources are equal (using the same mechanism as defined in the multiple codec case). Then, if the CMS wants to switch codecs it will need to place the new codec in the LocalConnectionOptions of a subsequent modify connection. However, there are certain risks with this approach. For example, when a modem call is detected and reported to the CMS, it may be possible that the resulting modify connection to use G.711 fails due to insufficient resources on the CMTS. This would not be the case if multiple codecs were defined since the LUB would already be reserved and guaranteed accessible for a subsequent commit.

5.6.11 V.152 support

In order to efficiently support the use of V.152 voice band data (VBD) in the DQoS network, an optimization to the MTA LUB behaviour described in clauses 5.6.10 and 6.1.1 is required. The following paragraph describes this enhancement.

When V.152 VBD is in the negotiated list of codecs, the MTA, when reserving the resources associated with the session, MUST reserve the resources for the negotiated codec list without taking into account the resources required for V.152 VBD. When the MTA detects a switch to V.152 VBD is required, the MTA must re-reserve and re-commit the resources required for V.152 VBD usage. The MTA should attempt this re-reservation and re-committal as soon as possible (after the discovery that a switch to V.152 VBD is required for the session), to meet the strict timing requirements of some analog modems. Once the bandwidth for V.152 VBD is reserved/committed, the full V.152 VBD bandwidth MUST remain reserved/committed for the duration of the call unless instructed otherwise via the CMS.

For V.152 VBD support, the CMS continues to authorize enough resources to satisfy the LUB of the session (including the resources needed for V.152 VBD). This allows the MTA to re-reserve and re-commit during the session without requiring authorization of additional resources from the CMS.

5.6.12 MTA port-to-port calls

When voice calls are established between different ports (endpoints) on the same MTA, DOCSIS forwarding rules specify that the CM must not forward packets over the DOCSIS network. As a result, the actions taken by the CMS and MTA in this special circumstance are different from a typical MTA-to-MTA call flow. A port-to-port call is defined by both endpoints using the same IP address.

If an MTA receives a connection request without a GateID, it MUST NOT initiate any DSx messaging to the CMTS. If an MTA is instructed to make a port-to-port call, the MTA MUST NOT initiate any DSx messaging to set up a service flow for this connection and MUST NOT send any of the voice packets over the network. In addition, if the MTA has previously created a service flow for a call where the far-end SDP was not available (but a GateID was specified in a CRCX or MDCX), then it MUST subsequently tear down the service flow if a port-to-port call is recognized once the remote SDP is received.

The CMS SHOULD recognize port-to-port calls and SHOULD omit Gate Control to the CMTS and SHOULD omit the GateID in the connection command to the MTA. Similar to the MTA case above, if the CMS has already established a gate for a call where the remote SDP was not available, it SHOULD expect a Gate-Close message from the CMTS once the MTA tears down the service flow upon detection of the port-to-port call. The CMS MUST NOT tear down a call between endpoints with the same IP address on receipt of a Gate-Close message.

5.6.13 Multiple grants per interval

For efficient use of DOCSIS resources, the MTA MAY choose to place multiple sub-flows with the same QoS parameter sets on the same service flow. Since the ServiceFlowScheduling type is part of the QoS parameter set, it MUST be common for all sub-flows that use the same DOCSIS Service Flow. For example, if a flow supporting silence suppression uses UGS/AD, and the existing Service Flow is configured for just UGS the new flow MUST be created on a separate service flow. For ease of implementation, when using multiple grants per interval, the existing service flow scheduling type cannot be changed.

The CMTS MAY support grants per interval greater than 1. If an MTA requests multiple grants per interval and the DSx message is rejected by the CMTS (i.e., the CMTS scheduler cannot properly schedule this request on the existing Service Flow, but may be able to service this request on a

separate service flow), the MTA MAY re-try using a separate service flow for the request (resources permitting).

The active grants per interval field in the extended MAC header is used to keep track of the active grants on a particular service flow that contains multiple sub-flows. For example, if you have two active calls and one enters silence suppression, the active grants in extended MAC header is decreased from 2 to 1. In this scenario, no DSC refresh is required on the flow because activity detection is flow based not grant based. The grants per interval in the DSC remains at 2 for Admitted and Active and flow refresh would only be required when the active grants goes to 0 all sub-flows enter silence suppression. The active grants per interval MUST be less than or equal to the number of sub-flows.

The PHS rules for all sub-flows in a service flow MUST be the same.

5.7 Theory of operation

5.7.1 Basic session set-up

Resource reservation is partitioned into separate Reserve and Commit phases. At the end of the first phase, resources are reserved but are not yet available to the MTA. (On the DOCSIS links, the service flows in each direction are admitted.) At the end of the second phase, resources are made available to the MTA and usage recording is started so that the user can be billed for usage. (On the DOCSIS links, the service flows are active.)




Figure 2 – Resource management phase 1

Figure 2 shows the first phase of the resource management protocol for a call. In this description, subscripts "O" and "T" designate the originating and terminating points of the call. As shown in Figure 2, MTA_O and MTA_T request resource reservation (DOCSIS Dynamic Services signalling for embedded clients) to CMTS_O and CMTS_T respectively. CMTS_O and CMTS_T perform an admission control check for resource availability (initiating signalling for resource reservation in the backbone if necessary) and send a reply to the respective MTAs which, in turn, respond to the CMS.

Figure 3 shows the second phase. After determining that resources are available, the CMS sends a message to MTA_T instructing it to start ringing the phone. When the called party picks up the phone, MTA_T sends a message to the CMS and the CMS instructs MTA_O and MTA_T to request a committal of resources. The arrival of the Commit messages at CMTS_T and CMTS_O causes them to open their gates, and also starts accounting for resource usage. To prevent some theft of service scenarios, each CMTS informs the respective CMS of the state change by sending a Gate-Open message.




Figure 3 – Resource management phase 2

5.7.2 Gate coordination

QoS signalling leads to the creation of a gate at each CMTS associated with a client involved in the session. Each gate maintains usage data for the session and controls whether the packets generated by the associated client receive access to enhanced QoS. Gate coordination is needed to prevent fraud and theft of service in situations where a malfunctioning or modified client does not issue the expected signalling messages. It is essential that protocol mechanisms are robust against abuse². A gate coordination protocol ensures that:

- A potential for one-way session establishment without billing is avoided. Because the clients may have adequate intelligence and are not trusted, one can envisage the clients establishing two one-way sessions to provide the users with an adequate interactive voice communication channel. Gate coordination prevents such sessions being established without the provider being able to charge for them.
- The opening and closing of each gate is closely synchronized with the corresponding state changes on the CMS.

5.7.3 Changing the packet classifiers associated with a gate

Once a pair of gates is set up, clients can communicate over the network with enhanced QoS. Several features needed for a commercial voice communications service involve changing the clients involved in a session, for example, when a session is transferred or redirected, or during three-way calling. This requires the packet classifiers associated with a gate to be modified to reflect the address of the new client. In addition, changing the endpoints involved in a session may affect how the session is billed. As a result, gates include addressing information for origination and termination points.

5.7.4 Session resources

The relationship between different categories of resources, authorized, reserved, and committed, is shown in Figure 4. A set of resources is represented by an n -dimensional space (shown here as two-dimensional) where n is the number of parameters (e.g., bandwidth, burst size, jitter, classifiers) needed to describe the resources. The exact procedures for comparing n -dimensional resource vectors are given in [ITU-T J.112].

When a session is first established, DQoS protocols authorize the use of some maximum amount of resources, indicated by the outer oval, specifying the authorized resources. When a client makes a reservation for a session, it reserves a certain amount of resources, which are not greater than those

² Several theft of service scenarios are described in Appendix IX.

for which it has been authorized. When the session is ready to proceed, the client commits to some amount of resources, which are not more than the reserved resources. In many common cases, the committed and reserved resources will be equal. The committed resources represent resources that are currently in use by the active session, whereas reserved resources represent those that are tied up by the client and have been removed from the pool for admission control purposes, but which are not necessarily being used by the client.




Figure 4 – Authorized, reserved and committed resources

Authorizations only affect future resource reservation requests. Resources that have been reserved prior to an authorization change are not affected.

Resources that have been reserved but not committed are available to the system for short-term uses only, such as handling of best-effort data. These resources are not available for other reservations (i.e., overbooking is not allowed). The maximum portion of the available resources that can be reserved at once is a policy decision by the CMTS, and outside the scope of DQoS.

Excess resources reserved above those committed are released unless the client explicitly requests they be kept through periodic reservation refresh operations. Maintaining such a condition for long periods of time is discouraged, as it reduces the overall capacity of the system. However, there are situations (e.g., call waiting service, where the call on hold requires resources beyond those needed for the active call) where excess reservations are necessary.

5.7.5 Admission control and session classes

It is envisaged that the Gate at the CMTS may use one or more session classes for resources reserved from an MTA. Session classes define provisioned admission control policies, or their parameters. It is expected that the provider would provision the necessary parameters and/or the alternative admission control policies in the CMTS and in the gate controller. For instance, a session class for normal voice communications, and an overlapping session class for emergency calls could be defined to allow the allocation of up to, respectively, 50% and 70% of the total resources to these classes of calls, and leaving the remainder 30-50% of the total bandwidth available to other, possibly lower priority, services. Session classes may furthermore enable pre-emption of already reserved resources, in which case the policy for such pre-emption would be provisioned by the service provider. When the Authorized Envelope is communicated to the Gate at the CMTS by the gate controller in the Gate-Set message, the gate controller includes adequate information to indicate which session class should apply when the corresponding DSA/DSC request is processed.

5.7.6 Resource renegotiations

Several of the supported session features require renegotiations of the QoS parameters associated with a session during the lifetime of the session. For example, clients might start communicating using a low-bit-rate audio codec. They can subsequently switch to a higher bit-rate codec or add a video stream, as long as the requested QoS is within the authorized envelope and there is available

bandwidth on the network. The use of an authorized QoS envelope that is pre-authorized by the gate controller acting as the policy decision point gives clients the flexibility to renegotiate QoS with the network without requiring subsequent gate controller involvement. This essentially means that use of resources up to the limits of the envelope is pre-authorized but NOT pre-reserved. Successful allocation of resources within the authorized envelope requires an admission control decision, and is not guaranteed. Subsequent to admission control, the resources are reserved for the flow, although the actual usage of the resources is permitted only after the Commit phase of the resource reservation protocol completes. However, no admission control decision is required at the time of commitment of resources. Each change in commitment of resources within the limits of the admission control decision does not require a further reservation. All reservation requests that pass admission control MUST fit within the authorization envelope.

5.7.7 Dynamic binding of resources (Re-reserve)

The dynamic QoS architecture recognizes that there may be a need to share resources across multiple sessions, especially when resources are in short supply. In particular, when using the call-waiting feature in telephony-like applications, the client may be involved in two simultaneous sessions, but will be active in only one conversation at a time. It is feasible in this case to share the network-layer resources (in particular, on the access link) between the two conversations. Therefore, this architecture allows a set of network layer resources (such as a bandwidth reservation) to be explicitly identified, and allows one or more gates to be associated with those resources. Signalling primitives allow the resources associated with a gate to be *shared* with another gate at the same CMTS. This improves the efficiency with which resources in the DOCSIS network are utilized.

When switching back and forth between two sessions in a call-waiting scenario, a client needs to keep enough resources reserved to accommodate either of the sessions, which in general may not need the same amount of resources. Thus, the re-commit operation may change the committed resources. However, the reserved resources do not change in this case, as the client should not have to go through admission control when switching back to the other session.

Whereas the committed resources are always associated with the current active session (and its corresponding IP flow), the reserved resources may be bound to different flows and different gates at different times. A handle, called a resource ID, is used to identify a set of reserved resources for the purpose of binding a flow to those resources.

5.7.8 Support for billing

QoS signalling can be used to support a broad range of billing models, based on only a stream of event records from the CMTS. Since the gate is in the data path, and since it participates in resource management interactions with a client, resource usage accounting is done by the gate. The gate in the CMTS is the appropriate place to do resource accounting, since the CMTS is directly involved in managing resources provided to a client. It is also important to do usage accounting in the CMTS to cope with client failures. If a client that is involved in an active session crashes, the CMTS MUST detect this and stop usage accounting for the session. This can be accomplished by monitoring the flow of packets along the data path for continuous-media applications, or by other mechanisms (such as station maintenance) performed by the CMTS. In addition, since the gate retains state for flows that have been authorized by a service-specific gate controller, it is used to hold service-specific information related to charging, such as the account number of the subscriber that will pay for the session. The policy function in the gate controller thus becomes stateless.

The support required in the CMTS is to generate and transmit an event message to a record keeping server on every change to the QoS, as authorized and specified by a gate. Opaque data provided by the gate controller that may be relevant to the record keeping server may also be included in the message. Requirements for handling of event records are contained in other Operations Support specifications.

5.7.9 Backbone resource management

When a CMTS receives a resource reservation message from an MTA, it first verifies that adequate upstream and downstream bandwidth is available over the access channel using locally available scheduling information. If this check is successful, the CMTS can either generate a new backbone resource reservation message, or forward towards the backbone a modified version of the resource reservation message received from the MTA. The CMTS performs any backbone-technology-specific mapping of the resource reservation that is needed. This enables the architecture to accommodate different backbone technologies, at the service provider's choosing. The specific mechanisms for reserving backbone QoS are outside the scope of this Recommendation.

A bidirectional model is used for resource reservation in the DOCSIS network where the routing is symmetric. A unidirectional model is used for resource reservation in the backbone, which allows routing asymmetries. Thus, when MTA_O makes a reservation with the CMTS, it knows two things: that it has adequate bandwidth in both directions over the DOCSIS network, and that it has adequate bandwidth over the backbone networks for the MTA_O to MTA_T flow. Thus, MTA_O knows that resources are available end-to-end in both directions once it gets a reply from MTA_T.

5.7.10 Setting the DiffServ code point

This architecture also allows for the use of a Differentiated Services backbone, where there is adequate bandwidth to carry voice conversations, but access to this bandwidth is on a controlled basis. Access to the bandwidth and differentiated treatment is provided to packets with the appropriate encoding of bits in the field of the IP header specified for Differentiated Service. This is called the DiffServ code point (DSCP). The DS field maintains backward compatibility with the present uses of the IP Precedence bits of the IPv4 TOS byte [IETF RFC 2474]. It is desirable to be able to set the DiffServ code point of packets that are about to enter the provider backbone from the CMTS. Since resources consumed by these packets in the backbone may depend heavily on this marking, this architecture provides control of the marking to network entities. This allows the network and service provider to control the use of the enhanced QoS rather than trusting the MTA. The provider can configure policies in the CMTS that determine how to set the DSCP for flows that pass through the CMTS. Such policies are sent to the CMTS in the gate set-up protocol from the CMS/GC.

For implementation efficiency, the information is passed to the MTA about the appropriate DSCP for it to use on a given session. The CMTS still needs to police received packets to ensure that correct DSCP is being used and that the volume of packets in a given class is within authorized bounds.

5.8 Sample mapping of SDP descriptions into RSVP flowspecs

Session descriptor protocol messages are used to describe multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation per [IETF RFC 2327]. This clause describes a mechanism for mapping the SDP description into RSVP flowSpecs.

A typical SDP description contains many fields that contain information regarding the session description (protocol version, session name, session attribute lines, etc.), the time description (time the session is active, etc.), and media description (media name and transport, media title, connection information, media attribute lines, etc.). The two critical components for mapping an SDP description into an RSVP FlowSpec message are the media name and transport address (m) and the media attribute lines (a).

The media name and transport address (m) are of the form:

m = <media> <port> <transport> <fmt list>

Copyright International Telecommunication Union/ITU Telecommunication Sector