

ETSI TS 101 909-4 V1.3.1 (2002-12)

Technical Specification

Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 4: Network Call Signalling Protocol



Reference

RTS/AT-020030-04

Keywords

access, broadband, cable, IP, multimedia, PSTN

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:

<http://www.etsi.org>

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at

<http://portal.etsi.org/tb/status/status.asp>

If you find errors in the present document, send your comment to:

editor@etsi.org

Copyright Notification

No part may be reproduced except as authorized by written permission.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2002.
All rights reserved.

DECT™, **PLUGTESTS™** and **UMTS™** are Trade Marks of ETSI registered for the benefit of its Members.
TIPHON™ and the **TIPHON logo** are Trade Marks currently being registered by ETSI for the benefit of its Members.
3GPP™ is a Trade Mark of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

Contents

Intellectual Property Rights	7
Foreword.....	7
Introduction	7
1 Scope	9
2 References	9
3 Definitions and abbreviations.....	11
3.1 Definitions	11
3.2 Abbreviations	11
4 Void.....	12
5 Overview	12
5.1 Relation with H.323 standards	13
5.2 Relation with IETF standards.....	14
6 Media Gateway Control Interface (MGCI).....	14
6.1 Model and naming conventions.....	14
6.1.1 Endpoint names	15
6.1.1.1 Embedded client endpoint names.....	16
6.1.1.1.1 Analogue access line endpoints.....	16
6.1.2 Call names	16
6.1.3 Connection names.....	17
6.1.4 Names of Call Agents and other entities.....	17
6.1.5 Digit maps.....	17
6.1.6 Events and signals.....	19
6.2 SDP use	20
6.3 Gateway control functions.....	21
6.3.1 NotificationRequest	22
6.3.2 Notifications	27
6.3.3 CreateConnection	28
6.3.4 ModifyConnection.....	32
6.3.5 DeleteConnection (from the Call Agent).....	33
6.3.6 DeleteConnection (from the Embedded Client).....	34
6.3.7 DeleteConnection (Multiple Connections From the Call Agent).....	35
6.3.8 Auditing	35
6.3.8.1 AuditEndPoint.....	36
6.3.8.2 AuditConnection	38
6.3.9 Restart in Progress	38
6.4 States, failover and race conditions	40
6.4.1 Recaps and highlights	40
6.4.2 Retransmission and detection of lost associations	40
6.4.3 Race conditions.....	43
6.4.3.1 Quarantine list.....	43
6.4.3.2 Explicit detection	46
6.4.3.3 Transactional semantics	46
6.4.3.4 Ordering of commands and treatment of disorder.....	47
6.4.3.5 Fighting the Restart Avalanche	47
6.4.3.6 Disconnected endpoints	48
6.5 Return codes and error codes	50
6.6 Reason codes.....	51
6.7 Use of local connection options and connection descriptors.....	51
7 Media Gateway Control Protocol.....	52
7.1 General description.....	52
7.2 Command header.....	53

7.2.1	Command line.....	53
7.2.1.1	Requested verb coding.....	53
7.2.1.2	Transaction identifiers.....	54
7.2.1.3	Endpoint, Call Agent and NotifiedEntity name coding.....	54
7.2.1.4	Protocol version coding.....	54
7.2.2	Parameter lines.....	55
7.2.2.1	Response acknowledgement.....	57
7.2.2.2	RequestIdentifier.....	57
7.2.2.3	Local connection options.....	57
7.2.2.4	Capabilities.....	58
7.2.2.5	Connection parameters.....	59
7.2.2.6	Reason codes.....	60
7.2.2.7	Connection mode.....	60
7.2.2.8	Event/signal name coding.....	60
7.2.2.9	RequestedEvents.....	61
7.2.2.10	SignalRequests.....	62
7.2.2.11	ObservedEvents.....	62
7.2.2.12	RequestedInfo.....	63
7.2.2.13	QuarantineHandling.....	63
7.2.2.14	DetectEvents.....	63
7.2.2.15	EventStates.....	63
7.2.2.16	ResourceID.....	63
7.2.2.17	RestartMethod.....	64
7.2.2.18	VersionSupported.....	64
7.3	Response header formats.....	64
7.3.1	CreateConnection.....	65
7.3.2	ModifyConnection.....	66
7.3.3	DeleteConnection.....	66
7.3.4	NotificationRequest.....	66
7.3.5	Notify.....	66
7.3.6	AuditEndpoint.....	67
7.3.7	AuditConnection.....	67
7.3.8	RestartInProgress.....	67
7.4	Session description encoding.....	67
7.4.1	SDP audio service use.....	68
7.4.1.1	Protocol version (v=).....	68
7.4.1.2	Origin (o=).....	68
7.4.1.3	Session name (s=).....	69
7.4.1.4	Session and media information (i=).....	69
7.4.1.5	URI (u=).....	69
7.4.1.6	E-mail address and phone number (e=, p=).....	69
7.4.1.7	Connection data (c=).....	69
7.4.1.8	Bandwidth (b=).....	70
7.4.1.9	Time, repeat times and time zones (t=, r=, z=).....	70
7.4.1.10	Encryption keys.....	70
7.4.1.11	Attributes (a=).....	71
7.4.1.12	Media Announcements (m=).....	73
7.4.2	SDP video service use.....	74
7.5	Transmission over UDP.....	74
7.5.1	Reliable message delivery.....	74
7.5.2	Retransmission strategy.....	74
7.6	Piggy-backing.....	75
7.7	Transaction identifiers and three ways handshake.....	75
7.8	Provisional responses.....	77
8	Security.....	77
Annex A (normative):	Event Packages.....	78
Annex B (normative):	Application of the NCS protocol to a SCN IPAT.....	84
B.1	Overview.....	84

B.2	Voice gateway architecture	84
B.3	Electrical and physical interface requirements	85
B.4	NCS package for V5 SCN protocol messages	87
B.4.1	Cadence-ringing request	87
B.4.1.1	Cadence ringing defaults and ranges	87
B.4.2	Pulsed signal request	88
B.4.2.1	Line treatment encoding	88
B.4.2.2	Line treatment defaults and ranges	89
B.4.2.3	Requested events	89
B.4.2.4	Pulse encoding	89
B.4.2.4.1	Pulse duration encoding	89
B.4.2.4.2	Pulse period encoding	89
B.4.2.5	Pulse completion event coding	90
B.4.2.6	Metering pulse report coding	90
B.4.2.7	V5 suppression indicator	90
B.4.2.7.1	No suppression	91
B.4.2.7.2	Suppression by pre-defined V5 signal message	91
B.4.2.7.3	Suppression by pre-defined line signal from TE	91
B.4.2.7.4	Suppression by pre-defined V5 SIGNAL message from LE or pre-defined line signal from TE	91
B.4.2.8	Repetition indicator	91
B.4.3	Pulse repetition encoding	92
B.4.4	Parameter usage	92
B.4.5	Pulsed signal cancellation	92
B.4.6	Pulsed completion event	92
B.4.7	Pulsed signal failure event	93
B.4.8	Steady-signal request	93
B.4.8.1	Line treatment encoding	93
B.4.8.2	Line treatment provisioning	93
B.4.9	Metering pulse generation	93
B.5	Provisioning configurations	94
B.5.1	MTA	94
B.5.2	IPAT	94
B.6	ETSI line package support	94
B.6.1	NCS audit	94
B.6.2	Unsupported signals - PICS declaration	94
B.7	Call flow examples	95
B.7.1	Cadence ringing	95
B.7.1.1	Cadence ringing call flow for basic ring cadence	95
B.7.1.2	Cadence ringing - Ring splash followed by a ring cadence	95
B.7.1.3	Cadence ringing - Ring splash followed by "on hook" data, then ring cadence	96
B.7.2	Pulsed signal request	97
B.7.2.1	Pulse signal request for one loop open pulse	97
B.7.2.2	Pulsed signal with start acknowledgement	98
B.7.2.3	Pulsed signal with completion acknowledgement	99
B.7.2.4	Pulsed signal with pulse acknowledgement	100
B.7.2.5	Pulsed signal - Meter pulse with pulse acknowledgement	101
B.7.2.6	Pulsed signal - Meter pulse with pulse acknowledgement with tariff change	102
B.7.3	Fixed meter pulse application, completed	103
B.7.4	Steady signal line treatment	104
B.7.4.1	Steady signal line treatment - Reverse polarity	104
Annex C (normative):	Dynamic Quality of Service	105
Annex D (informative):	Mode interactions	112
Annex E (informative):	Compatibility information	116
Annex F (informative):	Metering support for IPCablecom NCS	117

F.1	Objectives.....	117
F.2	Automatic metering package.....	117
F.2.1	Package name.....	117
F.2.2	Local connection options.....	117
F.2.3	Events and signals.....	117
F.2.3.1	Meter pulse burst signal.....	118
F.2.3.2	Enable metering signal.....	119
F.2.4	Properties.....	119
F.2.5	Statistics.....	119
F.2.6	Procedures.....	119
F.3	Use cases, Example call flows.....	120
F.3.1	Meter pulse while off hook.....	120
F.3.2	Meter pulse while on hook.....	120
F.3.3	Regular call charge.....	120
F.3.4	Call set-up charge.....	120
F.3.5	Mid-call tariff change.....	121
F.3.6	Mid-call add-on charge.....	121
F.3.7	End of call.....	121
F.3.8	Audit endpoint.....	122
F.4	Terms.....	122
Annex G (informative):	Bibliography.....	123
History.....		125

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<http://webapp.etsi.org/IPR/home.asp>).

All published ETSI deliverables shall include information which directs the reader to the above source of information.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Access and Terminals (AT).

The present document is part 4 of a multi-part deliverable covering Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services. Full details of the entire series can be found in part 1 [27].

The present document is part 4 of the series of ETSI deliverables and specifies a profile of an application programming interface, Media Gateway Controller Interface (MGCI), and a corresponding protocol, Media Gateway Control Protocol (MGCP), for controlling Voice over IP (VoIP) embedded clients from external call control elements. The MGCP assumes a call control architecture where the call control "intelligence" is outside the media gateway and is handled by external call control elements. The profile, as described in the present document, will be referred to as the Network-based Call Signalling (NCS) Protocol. It is based on the Media Gateway Control Protocol (MGCP) 1.0 IETF RFC 2705 [16], which is the result of a merging of the Simple Gateway Control Protocol, and the IP Device Control (IPDC) family of protocols, as well as additional ideas incorporated during the protocol development process.

Introduction

The cable industry in Europe and across other Global regions has already deployed broadband cable television Hybrid Fibre Coax (HFC) data networks running the Cable Modem Protocol. The cable industry is in the rapid stages of deploying IP Voice and other time critical multimedia services over these broadband cable television networks.

The cable industry has recognized the urgent need to develop ETSI Technical Specifications aimed at developing interoperable interface specifications and mechanisms for the delivery of end-to-end advanced real time IP multimedia time critical services over bi-directional broadband cable networks.

IPCablecom is a set of protocols and associated element functional requirements developed to deliver Quality of Service (QoS) enhanced secure IP multimedia time critical communications services using packetized data transmission technology to a consumer's home over the broadband cable television Hybrid Fibre/Coaxial (HFC) data network running the Cable Modem protocol. IPCablecom utilizes a network superstructure that overlays the two-way data-ready cable television network. While the initial service offerings in the IPCablecom product line are anticipated to be Packet Voice, the long-term project vision encompasses packet video and a large family of other packet-based services.

The cable industry is a global market and therefore the ETSI standards are developed to align with standards either already developed or under development in other regions. The ETSI Specifications are consistent with the CableLabs/PacketCable set of specifications as published by the SCTE. An agreement has been established between ETSI and SCTE in the US to ensure, where appropriate, that the release of PacketCable and IPCablecom set of specifications are aligned and to avoid unnecessary duplication. The set of IPCablecom ETSI specifications also refers to ITU-SG9 draft and published recommendations relating to IP Cable Communication.

The whole set of multi-part ETSI deliverables to which the present document belongs specify a Cable Communication Service for the delivery of IP Multimedia Time Critical Services over a HFC Broadband Cable Network to the consumers home cable telecom terminal. "IPCablecom" also refers to the ETSI working group program that shall define and develop these ETSI deliverables.

Many cable television operators are upgrading their facilities to provide two way capability and using this capability to provide high speed IP data services per ITU-T Recommendations J.83 [1] and J.112 [2]. These operators now want to expand the capability of this delivery platform to include a variety of time critical services. The present document is one of a series of documents required to achieve this goal. It provides a network based call signalling protocol necessary to establish connections.

1 Scope

The present document is one of a set of specifications belonging to IPCablecom. It specifies a profile of an application programming interface, Media Gateway Controller Interface (MGCI), and a corresponding protocol, Media Gateway Control Protocol (MGCP), for controlling Voice over IP (VoIP) embedded clients from external call control elements. The MGCP is based on a call control architecture, where the call control "intelligence" resides outside the gateways and is handled by external call control elements. The profile, as described in the present document, is referred to as the Network Call Signalling (NCS) Protocol.

Annex B of the present document specifies an application of the NCS protocol to an IPAT that is able to emulate an Access Node of an ETSI compliant local exchange. This annex specifies the mapping of a subset of the NCS protocol to V5.2.

Annex F of the present document specifies optional NCS components for the delivery of metering pulses.

The present set of documents specifies IPCablecom, a set of protocols and associated element functional requirements. These have been developed to deliver high Quality of Service (QoS), enhanced secure IP multimedia time critical communication services to a consumer's home over a cable television Hybrid Fibre/Coaxial (HFC) data network.

NOTE 1: IPCablecom defines a set of documents in the framework of a network superstructure that overlays the two-way data-ready cable television network, e.g. as specified within ES 201 488 [5] and ES 200 800 [4].

While the initial service offerings in the IPCablecom product line are anticipated to be Packet Voice and Packet Video, the long-term project vision encompasses a large family of packet-based services. This may require in the future, not only careful maintenance control but also an extension of the present set of documents.

NOTE 2: The present set of documents aims for global acceptance and applicability. It is therefore developed in alignment with standards either already existing or under development in other regions and in International Telecommunications Union (ITU).

NOTE 3: There may also be relevant equivalent studies on-going in IETF or ITU specific to the needs of European IPCablecom. Where possible ETSI TC AT-D working group for IPCablecom should align annex B package proposal or similar packages addressing the specific needs of European IPCablecom with work on going in IETF, Cablelabs, or the ITU.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication and/or edition number or version number) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

Referenced documents which are not found to be publicly available in the expected location might be found at <http://docbox.etsi.org/Reference>.

- [1] ITU-T Recommendation J.83: "Digital multi-programme systems for television, sound and data services for cable distribution".
- [2] ITU-T Recommendation J.112: "Transmission systems for interactive cable television services".
- [3] ETSI EN 300 324-1: "V interfaces at the digital Local Exchange (LE); V5.1 interface for the support of Access Network (AN); Part 1: V5.1 interface specification".
- [4] ETSI ES 200 800: "Digital Video Broadcasting (DVB); DVB interaction channel for Cable TV distribution systems (CATV)".

- [5] ETSI ES 201 488: "Data-Over-Cable Service Interface Specifications Radio Frequency Interface Specification".
- [6] ETSI TS 101 909-3: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 3: Audio Codec Requirements for the Provision of Bi-Directional Audio Service over Cable Television Networks using Cable Modems".
- [7] IETF RFC 821: "Simple Mail Transfer Protocol".
- [8] IETF RFC 1034: "Domain names - concepts and facilities".
- [9] IETF RFC 1889: "RTP: A Transport Protocol for Real-Time Applications".
- [10] IETF RFC 1890: "RTP Profile for Audio and Video Conferences with Minimal Control".
- [11] IETF RFC 2045: "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies".
- [12] IETF RFC 2234: "Augmented BNF for Syntax Specifications: ABNF".
- [13] IETF RFC 2326: "Real Time Streaming Protocol (RTSP)".
- [14] IETF RFC 2327: "SDP: Session Description Protocol".
- [15] IETF RFC 2543: "SIP: Session Initiation Protocol".
- [16] IETF RFC 2705: "Media Gateway Control Protocol (MGCP) Version 1.0".
- [17] ETSI EN 300 659-1: "Access and Terminals (AT); Analogue access to the Public Switched Telephone Network (PSTN); Subscriber line protocol over the local loop for display (and related) services; Part 1: On-hook data transmission".
- [18] ETSI EN 300 001: "Attachments to the Public Switched Telephone Network (PSTN); General technical requirements for equipment connected to an analogue subscriber interface in the PSTN".
- [19] EPC-RequDoc-V10-220501 (2001): "European Requirements for the Delivery of Time-critical Services over Cable Television Networks using IPcablecom".
- [20] ITU-T Recommendation G.711: "Pulse code modulation (PCM) of voice frequencies".
- [21] ETSI EG 201 188: "Public Switched Telephone Network (PSTN); Network Termination Point (NTP) analogue interface; Specification of physical and electrical characteristics at a 2-wire analogue presented NTP for short to medium length loop applications".
- [22] ETSI EN 300 659-3: "Access and Terminals (AT); Analogue access to the Public Switched Telephone Network (PSTN); Subscriber line protocol over the local loop for display (and related) services; Part 3: Data link message and parameter codings".
- [23] ITU-T Recommendation V.25: "Automatic answering equipment and general procedures for automatic calling equipment on the general switched telephone network including procedures for disabling of echo control devices for both manually and automatically established calls".
- [24] ETSI EN 300 347-1: "V interfaces at the digital Local Exchange (LE); V5.2 interface for the support of Access Network (AN); Part 1: V5.2 interface specification".
- [25] ETSI EN 300 166: "Transmission and Multiplexing (TM); Physical and electrical characteristics of hierarchical digital interfaces for equipment using the 2 048 kbit/s - based plesiochronous or synchronous digital hierarchies".
- [26] ETSI EN 300 167: "Transmission and Multiplexing (TM); Functional characteristics of 2 048 kbit/s interfaces".
- [27] ETSI TS 101 909-1: "Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 1: General".

- [28] ETSI TS 101 909-11: "Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 11: Security".
- [29] IETF RFC 2833: "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals".
- [30] ETSI ES 201 970: "Access and Terminals (AT); Public Switched Telephone Network (PSTN); Harmonized specification of physical and electrical characteristics at a 2-wire analogue presented Network Termination Point (NTP)".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

Access Node (AN): layer two termination device that terminates the network end of the ITU-T Recommendation J.112 connection

NOTE: It is technology specific. In ITU-T Recommendation J.112, annex A, it is called the INA while in annex B it is the CMTS.

cable modem: layer two termination device that terminates the customer end of the J.112 connection

embedded multimedia terminal adapter: physical and logical interface but it is physically integrated with the HFC Cable Modem functionality

NOTE: The difference in MTA and E-MTA defines product design architectures and is not relevant to this annex therefore these terms and their abbreviations are used synonymously and inter-changeably throughout the present document.

internet protocol access node: device in which the PacketCable Signalling Gateway (SG), Media Gateway (MG) and Media Gateway Controller (MGC) have been closely integrated

NOTE: The MGC includes the subset of Call Management Server (CMS) (aka Call Agent) functionality required to support the inter-working between a PacketCable HFC access network and a SCN through a switched circuit, line-controlled interface such as V5.

IPCablecom: ETSI working group project that includes an architecture and a series of specifications that enable the delivery of real time services (such as telephony) over the cable television networks using cable modems

line treatment: signals that can be applied to or sensed from the subscriber's loop, such as loop open, loop closed, ring, reverse battery, etc.

Multimedia Terminal Adapter (MTA): physical and logical interface between a line presence and the HFC Cable Modem (CM)

V5: general descriptor used to reference the V5.1/V5.2 Signalling protocols at the digital Local Exchange

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AAD	Average Acknowledgement Delay
ADEV	Average DEVIation
AN	Access Node
API	Application Programming Interface
AUCX	AuditConnection
AUEP	AUditEndPoint
BR	BRief signal
CA	Call Agents
CM	Cable Modem

CMS	Call Management Server
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CRCX	CReateConnection
DLCX	DeLeteConnection
DNS	Domain Name System
D-QoS	Dynamic Quality of Service
DTMF	Dual Tone Multi Frequency
E-MTA	Embedded Multimedia Terminal Adapter
HFC	Hybrid Fibre Coax
INA	Interactive Network Adaptor
IP	Internet Protocol
IPAT	Internet Protocol Access Terminal
IPDC	IP Device Control
LE	Local Exchange
MDCX	MoDifyConnection
MGCI	Media Gateway Controller Interface
MGCP	Media Gateway Control Protocol
MIB	Management Information Base
MTA	Media Terminal Adaptor
MWD	Maximum Waiting Delay
NCS	Network Call Signalling
NTFY	NoTiFY
OO	On/Off
PICS	Protocol Implementation Compliance Statement
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RQNT	NotificationRequest
RSIP	ReStartInProgress
RTCP	Real Time Control Protocol
RTP	Real-Time Protocol
RTSP	Real-Time Streaming Protocol
SAP	Session Announcement Protocol
SCN	Switched Circuit Network
SDP	Session Description Protocol
SIP	Session Initiation Protocol
SNMP	Simple Network Management Protocol
SSRC	Synchronization SouRCe
TDD	Telecomm Devices for the Deaf tones
TE	Terminal Equipment
TO	Time-Out
TTL	Time To Live
UDP	User Datagram Protocol
VoIP	Voice over IP

4 Void

5 Overview

The present document describes the NCS profile of an application programming interface (MGCI) and a corresponding protocol (MGCP) for controlling embedded clients from external call control elements. An embedded client is a network element that provides:

- two or more traditional analogue access lines to a Voice over IP (VoIP) network;
- one or more video lines to a VoIP network are for further study.

Embedded clients may not be confined to residential use only. For example, they may be used in a business as well. Embedded clients are used for line-side access and, as such, are expected to have line-side equipment, e.g. analogue access lines for conventional telephones associated with them, as opposed to trunk gateways.

The MGCP assumes a call control architecture where the call control "intelligence" is outside the gateways and handled by external call-control elements referred to as Call Agents. The MGCP assumes that these call-control elements, or Call Agents (CAs), will synchronize with each other to send coherent commands to the gateways under their control. The MGCP defined in the present document does not define a mechanism for synchronizing Call Agents, although future IP-Cablecom specifications may specify such mechanisms.

The MGCP assumes a connection model where the basic constructs are endpoints and connections. A gateway contains a collection of endpoints, which are sources, or sinks, of data and could be physical or virtual.

An example of a physical endpoint is an interface on a gateway that terminates an analogue POTS connection to a phone, key system, PBX, etc. A gateway that terminates residential POTS lines (to phones) is called a residential gateway, an embedded client or an MTA. Embedded clients may optionally support video as well.

An example of a virtual endpoint is an audio source in an audio-content server. Creation of physical endpoints requires hardware installation, while creation of virtual endpoints can be accomplished by software. However, the NCS profile of MGCP only addresses physical endpoints.

Connections are point-to-point. A point-to-point connection is an association between two endpoints with the purpose of transmitting data between these endpoints. Once this association is established for both endpoints, data transfer between these endpoints can take place. The association is established by creating the connection as two halves; one on the origination endpoint, and one on the terminating endpoint.

Call Agents instruct the gateways to create connections between endpoints and to detect certain events, e.g. off-hook, and generate certain signals, e.g. ringing. It is strictly up to the Call Agent to specify how and when connections are made, between which endpoints they are made, as well as what events and signals are to be detected and generated on the endpoints. The gateway, thereby, becomes a simple device, without any call state, that receives general instructions from the Call Agent without any need to know about or even understand the concept of calls, call states, features, or feature interactions. When new services are introduced, customer profiles changed, etc., the changes are transparent to the gateway. The Call Agents implement the changes and generate the appropriate new mix of instructions to the gateways for the changes made. Whenever the gateway reboots, it will come up in a clean state and simply carry out the Call Agent's instructions as they are received.

5.1 Relation with H.323 standards

The MGCP is designed as an internal protocol within a distributed system that appears to the outside as a single VoIP gateway. This system is composed of a Call Agent, which may or may not be distributed over several computer platforms, and a set of gateways. In an H.323 configuration, this distributed gateway system may interface on one side with one or more POTS lines, and on the other side with H.323 conformant systems, as illustrated in figure 1.




Figure 1

In the MGCP model, the gateways focus on the audio signal translation function, while the Call Agent handles the signalling and call processing functions. As a consequence, the Call Agent implements the "signalling" layers of the H.323 standard, and presents itself as an "H.323 Gatekeeper" or as one or more "H.323 Endpoints" to the H.323 systems. The H.225.0 call signalling and H.245 media signalling is therefore routed to the Call Agent.

5.2 Relation with IETF standards

While H.323 used to be the recognized standard for VoIP terminals, the IETF also has produced specifications for other types of multimedia applications. These other specifications include:

- the Session Description Protocol (SDP), IETF RFC 2327 [14];
- the Session Announcement Protocol (SAP), work in progress (see bibliography);
- the Session Initiation Protocol (SIP), IETF RFC 2543 [15];
- the Real-Time Streaming Protocol (RTSP), IETF RFC 2326 [13].

The latter three specifications are, in fact, alternative signalling standards that allow for the transmission of a session description to an interested party. SAP is used by multicast session managers to distribute a multicast session description to a large group of recipients. SIP is used to invite an individual user to take part in a point-to-point or unicast session. RTSP is used to interface a server that provides real-time data. In all three cases, the session description is described according to SDP; when audio is transmitted, it is transmitted through the real-time transport protocol (RTP and RTCP).

The distributed gateway systems and MGCP will enable PSTN voice communication and embedded client users to access sessions set up using SAP, SIP, or RTSP defined by the IETF MMUSIC Working Group. The Call Agent provides for signalling conversion, as illustrated in figure 2.

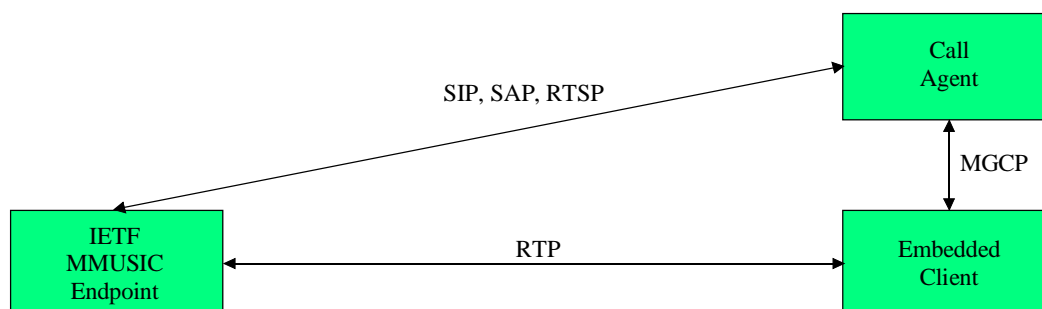


Figure 2

The SDP standard has a pivotal status in this architecture. We will see in the following description that we also use it to carry session descriptions in MGCP.

6 Media Gateway Control Interface (MGCI)

MGCI functions provide for connection control, endpoint control, auditing and status reporting. They each use the same system model and the same naming conventions.

6.1 Model and naming conventions

The MGCP assumes a connection model where the basic constructs are endpoints and connections. Connections are grouped in calls. One or more connections can belong to one call. Connections and calls are set up at the initiative of one or several Call Agents.

6.1.1 Endpoint names

Endpoint names, a.k.a. endpoint identifiers, have two components, both of which are defined to be case insensitive here:

- the domain name of the gateway managing the endpoint;
- a local endpoint name within that gateway.

Endpoint names will be of the form:

- `local-endpoint-name@domain-name`

where domain-name is an absolute domain-name as defined in IETF RFC 1034 [8] and includes a host portion, thus an example domain-name could be:

- `MyEmbeddedClient.cablelabs.com`

Also, domain-name may be an IPv4 address in dotted decimal form represented as a text-string and surrounded by a left and a right square bracket ("[" and "]") as in "[128.96.41.1]" - please consult IETF RFC 821 [7] for details. However, use of IP addresses is generally discouraged.

Embedded clients may have one or more endpoints (e.g. one for each RJ11 jack for black phones) associated with them, and each of the endpoints is identified by a separate local endpoint name. Just like the domain-name, the local endpoint name is case insensitive. Associated with the local endpoint name is an endpoint-type, which defines the type of the endpoint, such as analogue phone or video phone. The endpoint-type can be derived from the local endpoint name. The local endpoint name is a hierarchical name, where the least specific component of the name is the leftmost term, and the most specific component is the rightmost term. More formally, the local endpoint name must adhere to the following naming rules:

- The individual terms of the local endpoint name must be separated by a single slash ("/", ASCII 2F hex).
- The individual terms are ASCII character strings composed of letters, digits or other printable characters, with the exception of characters used as delimiters in endpoint-names ("/", "@"), characters used for wildcarding ("*", "\$"), and white space characters.
- Wild carding is represented either by an asterisk ("*") or a dollar sign ("\$") for the terms of the naming path which are to be wild-carded. Thus, if the full local endpoint name looks like:
`term1/term2/term3`
 and one of the terms of the local endpoint name is wild-carded, then the local endpoint name looks like this:
`term1/term2/*` if `term3` is wild-carded.
`term1/*/*` if `term2` and `term3` are wild-carded.
 In each of the examples, a dollar sign could have appeared instead of the asterisk.
- Wild-carding is only allowed from the right, thus if a term is wild-carded, then all terms to the right of that term must be wild-carded as well.
- In cases where mixed dollar sign and asterisk wild-cards are used, dollar-signs are only allowed from the right, thus if a term had a dollar sign wild-card, all terms to the right of that term must also contain dollar sign wild-cards.
- A term represented by an asterisk is to be interpreted as:
 - "use all values of this term known within the scope of the embedded client in question".
- A term represented by a dollar sign is to be interpreted as:
 - "use any one value of this term known within the scope of the embedded client in question".
- Each endpoint-type may specify additional detail in the naming rules for that endpoint-type, however such rules must not be in conflict with the above.

It should be noted that different endpoint-types or even different sub-terms, e.g. "lines", within the same endpoint-type will result in two different local endpoint names. Consequently, each "line" will be treated as a separate endpoint.

6.1.1.1 Embedded client endpoint names

Endpoints in embedded clients **MUST** support the additional naming conventions specified in this clause.

Embedded clients **MAY** support one or more endpoint-types including the following:

- Analogue Telephone: The analogue telephone is represented as an analogue access line (aaln). This is basically the equivalent of an analogue telephone line as known in the PSTN.
- Video: The details of the video device-type are for further study.
- Basic Access ISDN: The details of the ISDN device-type are for further study.

6.1.1.1.1 Analogue access line endpoints

In addition to the naming conventions specified above, local endpoint names for endpoints of type "analogue access line" (aaln) for embedded clients must adhere to the following:

- Local endpoint names contain at least one and, at most, two terms.
- Term1 **MUST** be the term "aaln" or a wildcard character. It should be noted that the use of a wildcard character for term1 could refer to any or all endpoint-types in the embedded client regardless of their type. Use of this feature is generally expected to be for administrative purposes, e.g. auditing or restart.
- Term2 **MUST** be a number from one to the number of analogue access lines supported by the embedded client in question. The number thus identifies a specific analogue access line on the embedded client.
- If a local endpoint name is composed of only one term, that term will be term1.
- If term1 *is not* a wildcard character, the wildcard character dollar sign (referring to "any one") is then assumed for term2, i.e. "aaln" is equivalent to "aaln/\$".
- If term1 *is* a wildcard character, the wildcard character asterisk (referring to "all") is then assumed for term2, i.e. "*" and "\$" is equivalent to respectively "*/*" and "\$/*".

Example analogue access line local endpoint names could thus be:

aaln/1	The first analogue access line on the embedded client in question.
aaln/2	The second analogue access line on the embedded client in question.
aaln/\$	Any analogue access line on the embedded client in question.
aaln/*	All analogue access lines on the embedded client in question.
*	All endpoints (regardless of endpoint-type) on the embedded client in question.

The provisioning/(auto)configuration process is responsible for obtaining and providing information about how many endpoints an embedded client has, as well as the endpoint-type of each endpoint. Although they are logically different, it should be noted that the *endpoint-type* can be derived from the local portion of the endpoint name.

6.1.2 Call names

Calls are identified by unique identifiers, independent of the underlying platforms or agents. Call identifiers are hexadecimal strings, which are created by the Call Agent. Call identifiers with a maximum length of 32 **MUST** be supported.

At a minimum, call identifiers **MUST** be unique within the collection of Call Agents that control the same gateways. However, the coordination of these call identifiers between Call Agents is outside the scope of the present document. When a Call Agent builds several connections that pertain to the same call, either on the same gateway or in different gateways, these connections all will be linked to the same call through the call identifier. This identifier then can be used by accounting or management procedures, which are outside the scope of MGCP.

6.1.3 Connection names

Connection identifiers are created by the gateway when it is requested to create a connection. They identify the connection within the context of an endpoint. Connection identifiers are treated in MGCP as hexadecimal strings. The gateway **MUST** ensure that a proper waiting period, at least three min, elapses between the end of a connection that used this identifier and its use in a new connection for the same endpoint. Connection names with a maximum length of 32 characters **MUST** be supported.

6.1.4 Names of Call Agents and other entities

The Media Gateway Control Protocol has been designed for enhanced network reliability to allow implementation of redundant Call Agents. This means that there is no fixed binding between entities and hardware platforms or network interfaces.

Call Agent names consist of two parts, similar to endpoint names. The local portion of the name does not exhibit any internal structure. An example Call Agent name is:

- `cal@ca.whatever.net`

Reliability is provided by the following precautions:

- Entities such as embedded clients or Call Agents are identified by their domain name, not their network addresses. Several addresses can be associated with a domain name. If a command cannot be forwarded to one of the network addresses, implementations **MUST** retry the transmission using another address.
- Entities may move to another platform. The association between a logical name (domain name) and the actual platform are kept in the Domain Name Service (DNS). Call Agents and gateways **MUST** keep track of the record's time-to-live read from the DNS. They **MUST** query the DNS to refresh the information if the time-to-live has expired.

In addition to the indirection provided by the use of domain names and the DNS, the concept of "notified entity" is central to reliability and fail-over in MGCP. The "notified entity" for an endpoint is the Call Agent currently controlling that endpoint. At any point in time, an endpoint has one, and only one, "notified entity" associated with it, and when the endpoint needs to send a command to the Call Agent, it **MUST** send the command to the current "notified entity" for which endpoint(s) the command pertains. Upon startup, the "notified entity" **MUST** be set to a provisioned value. Most commands sent by the Call Agent include the ability to explicitly name the "notified entity" through the use of a "NotifiedEntity" parameter. The "notified entity" **MUST** stay the same until either a new "NotifiedEntity" parameter is received or the endpoint reboots. If the "notified entity" for an endpoint is empty or has not been set explicitly (see note), the "notified entity" will then default to the source address of the last connection handling command or notification request received for the endpoint. Auditing will thus not change the "notified entity".

NOTE: This could happen as a result of specifying an empty NotifiedEntity parameter.

Clause 6.4 contains a more detailed description of reliability and fail-over.

6.1.5 Digit maps

The Call Agent can ask the gateway to collect digits dialled by the user. This facility is intended to be used for analogue access lines with residential gateways to collect the numbers that a user dials; it may also be used to collect access codes, credit card numbers, and other numbers requested by call control services. Endpoints **MUST** support digit maps as defined in this clause.

An alternative procedure involves the gateway notifying the Call Agent of the dialled digits as soon as they are dialled, a.k.a., overlap sending. However, such a procedure generates a large number of interactions. It is preferable to accumulate the dialled numbers in a buffer, and then to transmit them in a single message.

The problem with this accumulation approach, however, is that it is difficult for the gateway to predict how many numbers it needs to accumulate before transmission. For example, using the phone on our desk, we can dial the following numbers:

0	Local operator
00	Long distance operator
xxxx	Local extension number
8xxxxxxx	Local number
#xxxxxxx	Shortcut to local number at other corporate sites
*xx	Star services
91xxxxxxxxxx	Long distance number
9011 + up to 15 digits	International number

The solution to this problem is to load the gateway with a digit map that corresponds to the dial plan for the area in which the gateway resides. Thus the actual digit map used may differ between regions. This digit map is expressed using a syntax derived from the UNIX system command, *egrep*. For example, the dial plan described above results in the following digit map:

```
(0T| 00T|[1-7]xxx|8xxxxxxxx|#xxxxxxx|*xx|91xxxxxxxxxx|9011x.T)
```

The formal syntax of the digit map is described by the following BNF notation:

```
Digit ::= "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"
Timer ::= "T" | "t" -- matches the detection of a timer
Letter ::= Digit | Timer | "#" | "*" | "A" | "a" | "B" | "b" | "C" | "c" | "D" | "d"
Range ::= "X" | "x" -- matches any digit
| "[" Letters "]" -- matches any of the specified letters
Letters ::= Subrange | Subrange Letters
Subrange ::= Letter -- matches the specified letter

| Digit "-" Digit -- matches any digit between first and last
Position ::= Letter | Range
StringElement ::= Position -- matches an occurrence of the position
| Position "." -- matches an arbitrary number of occurrences
-- of the position, including 0
String ::= StringElement | StringElement String
StringList ::= String | String "|" StringList
DigitMap ::= String | "(" StringList ")"
```

A DigitMap, according to this syntax, is defined either by a (case insensitive) "string" or by a "list of strings" over which the gateway will attempt to find a shortest possible match.

Regardless of the above syntax, a timer is currently only allowed if it appears in the last position in a string (see note). Each string in the list is an alternate numbering scheme. A gateway that detects digits, letters, or timers will:

- 1) Add the event parameter code for the digit, letter, or timer, as a token to the end of the "current dial string" internal state variable.
- 2) Apply the "current dial string" to the digit map table, attempting a match to all expressions in the Digit Map.
- 3) If the result is under-qualified (partially matches at least one entry in the digit map and does not completely match another entry), do nothing further.

If the result matches an entry, or is over-qualified (i.e. no further digits could possibly produce a match), send the current dial string to the Call Agent (see 3)) and clear the "current dial string". A match, in the present document, can be either a "perfect match", exactly matching one of the specified alternative, or an impossible match, which occurs when the dial string does not match any of the alternatives. Unexpected timers, for example, can cause "impossible matches". Both perfect matches and impossible matches trigger notification of the accumulated digits (which may include other events).

Timer T is a digit input timer that can be used in two ways:

- When timer T is used with a digit map (see 3)), the timer is not started until the first digit is entered, and the timer is restarted after each new digit is entered until either a digit map match or mismatch occurs. In this case, timer T functions as an inter-digit timer.

NOTE: Technically speaking with the "accumulate according to digit map" action.

- When timer T is used without a digit map, the timer is started immediately and simply cancelled (but not restarted) as soon as a digit is entered. In this case, timer T can be used as an inter-digit timer when overlap sending is used.

When used with a digit map, timer T takes on one of two values, T_{par} or T_{crit} . When at least one more digit is required for the digit string to match any of the patterns in the digit map, timer T takes on the value T_{par} , corresponding to partial dial timing. If a timer is all that is required to produce a match, timer T takes on the value T_{crit} corresponding to critical timing. When timer T is used without a digit map, timer T takes on the value T_{crit} . The default value for T_{par} is 16 s and the default value for T_{crit} is 4 s. The provisioning process may alter both of these.

The end-points MUST support at least 2 048 bytes of digitmap; on all of the telephony interfaces.

Digit maps can be provided to the gateway by the Call Agent, whenever the Call Agent instructs the gateway to listen for digits. Again, it should be noted, that the details of the digit map used will depend on the area in which the gateway resides and thus the digit map is programmable. Digit maps, when provided by the Call Agent, MUST be as defined in this clause.

6.1.6 Events and signals

The concept of events and signals is central to MGCP. A Call Agent may ask to be notified about certain events occurring in an endpoint, e.g. off-hook events. A Call Agent also may request certain signals to be applied to an endpoint, e.g. dial-tone.

Events and signals are grouped in packages within which they share the same namespace, which we will refer to as event names in the following. A package is a collection of events and signals supported by a particular endpoint-type. For instance, one package may support a certain group of events and signals for analogue access lines, and another package may support another group of events and signals for video lines. One or more packages may exist for a given endpoint-type, and each endpoint-type has a default package with which it is associated.

Event names consist of a package name and an event code and, since each package defines a separate namespace, the same event codes may be used in different packages. Package names and event codes are case insensitive strings of letters, digits, and hyphens, with the restriction that hyphens MUST NOT be the first or last character in a name. Some event codes may need to be parameterized with additional data, which is accomplished by adding the parameters between a set of parentheses. The package name is separated from the event code by a slash ("/"). The package name may be excluded from the event name, in which case the default package name for the endpoint-type in question is assumed. For example, for an analogue access line with the example line package (package name "X") being the default package, the following two event names are considered equal:

- X/dl dial-tone in the example line package for an analogue access line;
- dl dial-tone in the example line package (default) for an analogue access line.

Annex A defines an initial set of packages. Additional package names and event codes may be defined by and/or registered with IPCablecom. Any change to the packages defined in the present document MUST result in a change of the package name, or a change in the NCS profile version number, or possibly both.

Each package MUST have a package definition, which MUST define the name of the package, and the definition of each event belonging to the package. The event definition MUST include the precise name of the event, i.e. the event code, a plain text definition of the event and, when appropriate, the precise definition of the corresponding signals, for example the exact frequencies of audio signals such as dial-tone or DTMF tones. Events must further specify if they are persistent (e.g. off-hook, see clause 6.3.1) and if they contain auditable event-states (e.g. off-hook, see clause 6.3.8.1). Signals MUST also have their type defined (On/Off, Time-Out, or Brief), and Time-Out signals MUST have a default time-out value defined - see clause 6.3.1.

In addition to IPCablecom packages, implementers MAY gain experience by defining experimental packages. The package name of experimental packages MUST begin with the two characters "x-" or "X-"; IPCablecom MUST NOT register package names that start with these two characters. An embedded client that receives a command referring to an unsupported package MUST return an error (error code 518 - unsupported package).

Package names and event codes support one wild-card notation each. The wildcard character "*" (asterisk) can be used to refer to all packages supported by the endpoint in question, and the event code "all" to refer to all events in the package in question. For example:

- X/all refers to all events in the example line package for an analogue access line;
- */all for an analogue access line; refers to all packages and all events in those packages supported by the endpoint in question.

Consequently, the package name "*" MUST NOT be assigned to a package, and the event code "all" MUST NOT be used in any package.

Events and signals are by default detected and generated on endpoints, however some events and signals may be detected and generated on connections in addition to or instead of on an endpoint. For example, endpoints may be asked to provide a ringback tone on a connection. In order for an event or signal to be able to be detected or generated on a connection, the definition of the event/signal MUST explicitly define that the event/signal can be detected or generated on a connection.

When a signal shall be applied on a connection, the name of the connection is added to the name of the event, using an "at" sign (@) as a delimiter, as in:

- X/rt@0A3F58

The wildcard character "*" (asterisk) can be used to denote "all connections" on the affected endpoint(s). When this convention is used, the gateway MUST generate or detect the event on all the connections that are connected to the endpoint(s). An example of this convention is:

- X/rt@*

The wildcard character "\$" (dollar sign) can be used to denote "the current connection". This convention MUST NOT be used unless the event notification request is "encapsulated" within a CreateConnection or ModifyConnection command. When the convention is used, the gateway MUST generate or detect the event on the connection that is currently being created or modified. An example of this convention is:

- X/rt@\$

The connection id, or a wildcard replacement, can be used in conjunction with the "all packages" and "all events" conventions. For example, the notation:

- */all@*

can be used to designate all events on all connections for the affected endpoint(s).

6.2 SDP use

The Call Agent uses the MGCP to provide the gateways with the description of connection parameters such as IP addresses, UDP port, and RTP profiles. Except where otherwise noted or implied in the present document, SDP descriptions MUST follow the conventions delineated in the Session Description Protocol (SDP), which is now an IETF-proposed standard IETF RFC 2327 [14].

SDP allows for description of multimedia conferences. The NCS profile will only support the setting of audio and video connections using the media types "audio" and "video". Currently, only "audio" connections have been specified.

6.3 Gateway control functions

This clause describes the commands of the MGCP in the form of a remote procedure call (RPC) like API, which we will refer to as the media gateway control interface (MGCI). An MGCI function is defined for each MGCP command, where the MGCI function takes and returns the same parameters as the corresponding MGCP command. The functions shown in this clause provide a high-level description of the operation of MGCP and describe an *example* of an RPC-like API that MAY be used for an implementation of MGCP. Although the MGCI API is merely an example API, the semantic behaviour defined by MGCI is an integral part of the specification, and all implementations MUST conform to the semantics specified for MGCI. The actual MGCP messages exchanged, including the message formats and encodings used are defined in the protocol section (see clause 15). Embedded clients MUST implement those exactly as specified.

The MGCI service consists of connection handling and endpoint handling commands. The following is an overview of the commands:

- The Call Agent can issue a NotificationRequest command to a gateway, instructing the gateway to watch for specific events such as hook actions or DTMF tones on a specified endpoint.
- The gateway will then use the Notify command to inform the Call Agent when the requested events occur on the specified endpoint.
- The Call Agent can use the CreateConnection command to create a connection that terminates in an endpoint inside the gateway.
- The Call Agent can use the ModifyConnection command to change the parameters associated to a previously established connection.
- The Call Agent can use the DeleteConnection command to delete an existing connection. In some circumstances, the DeleteConnection command also can be used by a gateway to indicate that a connection can no longer be sustained.
- The Call Agent can use the AuditEndpoint and AuditConnection commands to audit the status of an "endpoint" and any connections associated with it. Network management beyond the capabilities provided by these commands are generally desirable, e.g. information about the status of the embedded client. Such capabilities are expected to be supported by the use of the Simple Network Management Protocol (SNMP) and definition of a MIB, which is outside the scope of the present document.
- The gateway can use the RestartInProgress command to notify the Call Agent that the endpoint, or a group of endpoints managed by the gateway, is being taken out of service or is being placed back in service.

These services allow a controller (normally the Call Agent) to instruct a gateway on the creation of connections that terminate in an endpoint attached to the gateway, and to be informed about events occurring at the endpoint. Currently, an endpoint is limited to a specific analogue access line within an embedded client.

Connections are grouped into "calls". Several connections, that may or may not belong to the same call, can terminate in the same endpoint. Each connection is qualified by a "mode" parameter, which can be set to "send only" (sendonly), "receive only" (recvonly), "send/receive" (sendrecv), "conference" (confrnce), "inactive" (inactive), "replicate" (replcate), "network loopback" (netwloop) or "network continuity test" (netwtst). The "mode" parameter determines if media packets can be sent and/or received on the connection; however, RTCP is unaffected.

Audio signals received from the endpoint will be sent on any connection for that endpoint whose mode is either "send only", "send/receive", "conference", or "replicate".

Handling of the audio signals received on these connections is also determined by the mode parameters:

- Audio signals received in data packets through connections in "inactive" or "replicate" mode are discarded.
- Audio signals received in data packets through connections in "receive only", "conference", or "send/receive" mode are mixed together and then sent to the endpoint.
- Audio signals originating from the endpoint are transmitted over all the connections whose mode is "send only", "conference", or "send/receive".

- In addition to being sent to the endpoint, audio signals received in data packets through connections in "conference" mode are replicated to all the other connections for the endpoint whose mode is "conference". The details of this forwarding, e.g. RTP translator or mixer, etc., is outside the scope of the present document.
- Audio signals sent to and from the endpoint are mixed and transmitted over all the connections whose mode is "replicate". This SHOULD include audio signals generated by signals.
- Audio signals received in data packets through connections in "network loopback" or "network continuity test" mode will be sent back on the connection as described below.

If the mode is set to "network loopback," the audio signals received from the connection will be echoed back on the same connection. The "network loopback" mode SHOULD simply operate as an RTP packet reflector.

The "network continuity test" mode is used for continuity checking across the IP network. An endpoint-type specific signal is sent to the endpoints over the IP network, and the endpoint is then supposed to echo the signal over the IP network after passing it through the gateway's internal equipment to verify proper operation. The signal MUST go through internal decoding and re-encoding prior to being passed back. For analogue access lines, the signal will be an audio signal, and the signal MUST NOT be passed on to a telephone connected to the analogue access line, regardless of the current hook-state of that handset, i.e. on-hook or off-hook.

New and existing connections for the endpoint MUST NOT be affected by connections placed in "network loopback" or "network continuity test" mode. However, local resource constraints may limit the number of new connections that can be made.

The "replicate" mode MUST at a minimum support replicating the stream from the endpoint and one other connection regardless of the encoding method used for that other connection. The "replicate" connection is however only REQUIRED to support a resulting media stream in ITU-T Recommendation G.711 [20] encoding (see note). Support of the "conference" mode is optional. Please refer to annex F for illustrations of mode interactions.

NOTE: The "replicate" connection can, e.g. be used to support "busy line verification" with minimal resource impact on the embedded client.

6.3.1 NotificationRequest

The NotificationRequest command is used to request the gateway to send a notification upon the occurrence of specified events in an endpoint. For example, a notification may be requested when tones associated with fax communication are detected on the endpoint. The entity receiving this notification, usually the Call Agent, may then decide that a different type of encoding should be used on the connections bound to this endpoint and instruct the gateway accordingly (see note 1).

NOTE 1: The new instruction would be a ModifyConnection command.

```

ReturnCode
  ← NotificationRequest(EndpointId
    [, NotifiedEntity]
    [, RequestedEvents]
    [, RequestIdentifier]
    [, DigitMap]
    [, SignalRequests]
    [, QuarantineHandling]
    [, DetectEvents])

```

EndpointId is the identifier for the endpoint(s) in the gateway where NotificationRequest executes. The EndpointId MUST follow the rules for endpoint names specified in clause 6.1.1. The "any of" wildcard MUST NOT be used.

NotifiedEntity is an optional parameter that specifies a new "notified entity" for the endpoint.

RequestIdentifier is used to correlate this request with the notification it may trigger. It will be repeated in the corresponding Notify command.

SignalRequests is a parameter that contains the set of signals that the gateway is asked to apply. Unless otherwise specified, signals are applied to the endpoint, however some signals can be applied to a connection. The following are examples of signals (see note 2):

NOTE 2: Please refer to 0 for a complete list of signals.

- Ringing;
- Busy tone;
- Call waiting tone;
- Off hook warning tone;
- Ringback tones on a connection.

Signals are divided into different types depending upon their behaviour:

- **On/Off (OO)**: Once applied, these signals last until they are turned off. This can only happen as the result of a new SignalRequests where the signal is turned off (see later). Signals of type OO are defined to be idempotent, thus multiple requests to turn a given OO signal on (or off) are perfectly valid and **MUST NOT** result in any errors. An On/Off signal could be a visual message waiting indicator (VMWI). Once turned on, it **MUST NOT** be turned off until explicitly instructed to by the Call Agent, or the endpoint restarts.
- **Time-Out (TO)**: Once applied, these signals last until they are either cancelled (by the occurrence of an event or by not being included in a subsequent [possibly empty] list of signals), or a signal-specific period of time has elapsed. A signal that times out will generate an "operation complete" event (please see annex A for further definition of this event). A TO signal could be "ringback" timing out after 180 s. If an event occurs prior to the 180 s, the signal will, by default, be stopped (see note 3). If the signal is not stopped, the signal will time out, stop and generate an "operation complete" event, about which the Call Agent may or may not have requested to be notified. If the Call Agent has asked for the "operation complete" event to be notified, the "operation complete" event sent to the Call Agent will include the name(s) of the signal(s) that timed out (see note 4). Signal(s) generated on a connection will include the name of that connection. Time-out signals have a default time-out value defined for them, which may be altered by the provisioning process. Also, the time-out period may be provided as a parameter to the signal. A value of zero indicates that the time-out period is infinite. A TO signal that fails after being started, but before having generated an "operation complete" event will generate an "operation failure" event, which will include the name(s) of the signal(s), that timed out (see note 5).

NOTE 3: The "Keep signal(s) active" action may override this behaviour.

NOTE 4: If parameters were passed to the signal, the parameters will not be reported.

NOTE 5: These are merely examples from the example line package in 0.

- **Brief (BR)**: The duration of these signals is so short that they stop on their own. If a signal stopping event occurs, or a new SignalRequests is applied, a currently active BR signal will not stop. However, any pending BR signals not yet applied will be cancelled. A brief tone could be a DTMF digit. If the DTMF digit "1" is currently being played, and a signal stopping event occurs, the "1" would finish playing.

Signals are, by default, applied to endpoints. If a signal applied to an endpoint results in the generation of a media stream (audio, video, etc.), the media stream **MUST NOT** be forwarded on any connection associated with that endpoint, regardless of the mode of the connection. For example, if a call-waiting tone is applied to an endpoint involved in an active call, only the party using the endpoint in question will hear the call-waiting tone. However, individual signals may define a different behaviour.

When a signal is applied to a connection that has received a RemoteConnectionDescriptor (see clause 6.3.3), the media stream generated by that signal **MUST** be forwarded on the connection *regardless* of the current mode of the connection. If a RemoteConnectionDescriptor has not been received, the gateway **MUST** return an error (error code 527 - missing RemoteConnectionDescriptor).

When a (possibly empty) list of signal(s) is supplied, this list completely replaces the current list of active time-out signals. Currently active time-out signals that are not provided in the new list **MUST** be stopped and the new signal(s) provided will now become active. Currently active time-out signals that are provided in the new list of signals **MUST** remain active without interruption, thus the timer for such time-out signals will not be affected. Consequently, there is currently no way to restart the timer for a currently active time-out signal without turning the signal off first. If the time-out signal is parameterized, the original set of parameters **MUST** remain in effect, regardless of what values are provided subsequently. A given signal **MUST NOT** appear more than once in a SignalRequests.

The currently defined signals can be found in annex A.

RequestedEvents is a list of events that the gateway is requested to detect on the endpoint. Unless otherwise specified, events are detected on the endpoint, however some events can be detected on a connection. Examples of events are (see note 6):

- on-hook transition (occurring in classic telephone sets when the user hangs up the handset);
- off-hook transition (occurring in classic telephone sets when the user lifts the handset);
- DTMF digits (or pulse digits).

The currently defined events can be found in annex A.

NOTE 6: These are merely examples from the example line package in 0.

To each event is associated one or more **actions** that define the action that the gateway must take when the event in question occurs. The possible actions are:

- Notify the event immediately, together with the accumulated list of observed events.
- Accumulate the event.
- Accumulate according to Digit Map.
- Ignore the event.
- Keep Signal(s) active.
- Embedded NotificationRequest.
- Embedded ModifyConnection.

Two sets of requested events will be detected by the endpoint; persistent and non-persistent.

Persistent events are always detected on an endpoint. If a persistent event is not included in the list of RequestedEvents, and the event occurs, the event will be detected anyway, and processed like all other events, as if the persistent event had been requested with a Notify action (see note 7). Thus, informally, persistent events can be viewed as always being implicitly included in the list of RequestedEvents with an action to Notify, although no glare detection, etc., will be performed (see note 8). Persistent events are identified as such through their definition - see annex A.

NOTE 7: Thus the RequestIdentifier will be the RequestIdentifier of the current NotificationRequest.

NOTE 8: Normally, if a request to look for, e.g. off-hook, is made, the request is only successful if the phone is not already off-hook.

Non-persistent events are those events that have to be explicitly included in the RequestedEvents list. The (possibly empty) list of requested events completely replaces the previous list of requested events. In addition to the persistent events, only the events specified in the requested events list will be detected by the endpoint. If a persistent event is included in the RequestedEvents list, the action specified will then replace the default action associated with the event for the life of the RequestedEvents list, after which the default action is restored. For example, if "Ignore off-hook" was specified, and a new request without any off-hook instructions were received, the default "Notify off-hook" operation then would be restored. A given event **MUST NOT** appear more than once in a RequestedEvents.

More than one action can be specified for an event, although a given action can not appear more than once for a given event. The following matrix specifies the legal combinations of actions.

Table 1

	Notify	Accumulate	Accumulate according to digit map	Ignore	Keep Signal(s) Active	Embedded NotificationRequest	Embedded ModifyConnection
Notify	–	–	–	–	√	√	√
Accumulate	–	–	–	–	√	√	√
Accumulate according to digit map	–	–	–	–	√	–	√
Ignore	–	–	–	–	√	–	√
Keep Signal(s) active	√	√	√	√	–	√	√
Embedded NotificationRequest	√	√	–	–	√	–	√
Embedded ModifyConnection	√	√	√	√	√	√	–

If a client receives a request with an invalid action or illegal combination of actions, it MUST return an error to the Call Agent (error code 523 - unknown or illegal combination of actions).

When multiple actions are specified, e.g. "Keep signal(s) active" and "Notify", the individual actions are assumed to occur simultaneously.

The Call Agent can send a NotificationRequest with an empty RequestedEvents list to the gateway. The Call Agent can do so, for example, to an embedded client when it does not want to collect any more DTMF digits. However, persistent events will still be detected and notified.

DigitMap is an optional parameter that allows the Call Agent to provision the endpoint with a digit map according to which digits will be accumulated when the Call Agent provides a RequestedEvents parameter with the action "accumulate according to digit map" for that endpoint. The digit map provided is persistent and, therefore, need not be provided whenever a request to "accumulate according to digit map" is made, however Call Agents can provide a digit map at any time. A digit map MUST be provided for the endpoint no later than with the first request to "accumulate according to digit map". If the gateway is requested to "accumulate according to digit map" and the gateway currently does not have a digit map for the endpoint in question, the gateway MUST return an error (error code 519 - endpoint does not have a digit map).

Each endpoint has a variable called the "current dial string" in which digits are collected for matching with the digit map, as specified in clause 6.1.5. Whenever a Notify is sent or a NotificationRequest is to be processed, the "current dial string" is initialized to a null string. The digits to be processed may now either be detected as input, or they may be retrieved from an event input holding area known as the "quarantine buffer" - please see clause 6.4.3.1 for further details.

The signals being applied by the SignalRequests are synchronized with the collection of events specified or implied in the RequestedEvents parameter, except if overridden by the "Keep signal(s) active" action. For example, if the NotificationRequest mandated a "ringing" signal and the event request asked to look for an "off-hook" event, the ringing should, by default, stop as soon as the gateway detected an off-hook event. If "off-hook" was defined as a persistent event and the event request did not ask to look for an "off-hook" event, the ringing would stop anyway since off-hook would then be implied in the RequestedEvents parameter. The formal definition is that the generation of all "Time Out" signals MUST stop as soon as one of the requested events is detected, unless the "Keep signal(s) active" action is associated to the specified event. In the case of the action "accumulate according to digit map", the default behaviour would be to stop all active time-out signals when the first digit (see note 9) is accumulated - it is irrelevant to this synchronization if the accumulated digit results in a match, mismatch, or partial matching to the digit map.

NOTE 9: Digit as defined in digit maps, i.e. including asterisk, timer, etc.

If it is desired that time-out signal(s) continue when a looked-for event occurs, the "Keep Signal(s) Active" action can be used. This action has the effect of keeping all currently active time-out signal(s) active, thereby negating the default stopping of time-out signals upon the event's occurrence.

If signal(s) are desired to start when a looked-for event occurs, the "Embedded NotificationRequest" action can be used. The embedded NotificationRequest may include a new list of RequestedEvents, SignalRequests and a new Digit Map as well. The semantics of the embedded NotificationRequest is as if a new NotificationRequest was just received with the same NotifiedEntity, RequestIdentifier, QuarantineHandling and DetectEvents. When the "Embedded NotificationRequest" is activated, the "current dial string" will be cleared; however the list of observed events and the quarantine buffer will be unaffected (if combined with a Notify, the Notify will clear the ObservedEvents list though - see clause 6.4.3.1). Note, that the Embedded NotificationRequest action does not accumulate the triggering event, however it can be combined with the Accumulate action to achieve that. NCS implementations MUST be able to support at least one level of embedding. An embedded NotificationRequest that respects this limitation MUST NOT contain another Embedded NotificationRequest.

The embedded NotificationRequest action allows the Call Agent to set up a "mini-script" to be processed by the gateway immediately following the detection of the associated event. Any SignalRequests specified in the embedded NotificationRequest will start immediately. Considerable care must be taken to prevent discrepancies between the Call Agent and the gateway. However, long-term discrepancies should not occur as new SignalRequests completely replaces the old list of active time-out signals, and BR-type signals always stop on their own. Limiting the number of On/Off-type signals is encouraged. It is considered good practice for a Call Agent to occasionally turn on all On/Off signals that should be on, and turn off all On/Off signals that should be off.

If connection modes are desired to be changed when a looked-for event occurs, the "Embedded ModifyConnection" action can be used. The embedded ModifyConnection may include a list of connection mode changes each consisting of the mode change and the affected connection-id. The wildcard "\$" can be used to denote "the current connection", however this notation MUST NOT be used outside a connection handling command - the wildcard refers to the connection in question for the connection handling command.

The embedded ModifyConnection action allows the Call Agent to instruct the endpoint to change the connection mode of one or more connections immediately following the detection of the associated event. Each of connection mode changes work similarly to a corresponding ModifyConnection command (see note 10). When a list of connection mode changes is supplied, the connection mode changes MUST be applied one at a time in left-to-right order. When all the connection mode changes have finished, an "operation complete" event parameterized with the name of the completed action will be generated (see annex A for details). Should any of the connection mode changes fail, an "operation failure" event parameterized with the name of the failed action and connection mode change will be generated (see annex A for details) - the rest of the connection mode changes MUST NOT be attempted, and the previous successful connection mode changes in the list MUST remain effective.

NOTE 10: Thus, if, e.g. D-QoS is used on the connection, the default D-QoS action will still be taken when the embedded ModifyConnection action is carried out.

Finally, the Ignore action can be used to ignore an event, e.g. to prevent a persistent event from being notified. However, the synchronization between the event and an active signal will still occur by default.

Clause 6.4.3.1 contains additional details on the semantics of event detection and reporting. The reader is encouraged to study it carefully.

The specific definition of actions that are requested via these SignalRequests (e.g. the duration of and frequency of a DTMF digit) is outside the scope of the core NCS Specification. This definition may vary from location to location and, hence, from gateway to gateway. Consequently, the definitions are provided in event packages, which may be provided outside of the core specification. An initial list of event packages can be found in annex A.

The RequestedEvents and SignalRequests generally refer to the same events. In one case, the gateway is asked to detect the occurrence of the event and, in the other case, it is asked to generate it. There are exceptions to this rule, for example, fax and modem tones, which can be detected but can not be signalled. However, we necessarily cannot expect all endpoints to detect all events. The specific events and signals that a given endpoint can detect or perform are determined by the list of event packages that are supported by that endpoint. Each package specifies a list of events and signals that can be detected or applied. A gateway that is requested to detect or to apply an event that is not supported by the specified endpoint MUST return an error (error code 512 or 513 - not equipped to detect event or generate signal). When the event name is not qualified by a package name, the default package name for the endpoint is assumed. If the event name is not registered in this default package, the gateway MUST return an error (error code 522 - no such event or signal).

The Call Agent can send a NotificationRequest whose requested signal list is empty. This has the effect of stopping all active time-out signals. It can do so, for example, when tone generation, e.g. ringback, should stop.

QuarantineHandling is an optional parameter that specifies the handling options for events in the quarantine buffer (see clause 6.4.3.1), i.e. events that have been detected by the gateway before the arrival of this **NotificationRequest** command, but have not yet been notified to the Call Agent. The parameter provides a set of handling options: whether the quarantined events should be processed or discarded (the default is to process them.) whether the gateway is expected to generate at most one notification (lockstep), or multiple notifications (loop), in response to this request (the default is at most one). When the parameter is absent, the quarantined events **MUST** be processed. Support for the "lockstep" mode (via default) is mandatory while support for "loop" mode is optional. An endpoint that receives a **NotificationRequest** with an unsupported **QuarantineHandling** parameter value **SHOULD** respond with error code 508 (unsupported **QuarantineHandling**). Note that the quarantine-handling parameter also governs the handling of events that were detected and processed but not yet notified when the command is received.

DetectEvents is an optional parameter that specifies a minimum list of events that the gateway is requested to detect in the "notification" and "lockstep" state. When this parameter is absent, the events that **MUST** be detected in the quarantine period are those listed in the last received **DetectEvents** list. In addition, the gateway **MUST** also detect persistent events and the events specified in the **RequestedEvents** list, including those for which the "ignore" action is specified. Further explanation of this parameter may be found in clause 6.4.3.1.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

6.3.2 Notifications

Notifications are sent via the **Notify** command by the gateway when an observed event is to be notified:

```
ReturnCode
  ← Notify(EndpointId
           [, NotifiedEntity]
           , RequestIdentifier
           , ObservedEvents)
```

EndpointId is the name for the endpoint in the gateway, which is issuing the **Notify** command, as defined in clause 6.1.1. The identifier **MUST** be a fully qualified endpoint name, including the domain name of the gateway. The local part of the name **MUST NOT** use the wildcard convention.

NotifiedEntity is an optional parameter that identifies the entity to which the notification is sent. This parameter is equal to the **NotifiedEntity** parameter of the **NotificationRequest** that triggered this notification. The parameter is absent if there was no such parameter in the triggering request. Regardless of the value of the "NotifiedEntity" parameter, the notification **MUST** be sent to the current "notified entity" for the endpoint.

RequestIdentifier is a parameter that repeats the **RequestIdentifier** parameter of the **NotificationRequest** that triggered this notification. It is used to correlate this notification with the notification request that triggered it. Persistent events will be viewed here as if they had been included in the last **NotificationRequest**. When no **NotificationRequest** has been received, the **RequestIdentifier** used will be zero ("0").

ObservedEvents is a list of events that the gateway detected and accumulated, either by the "accumulate", "accumulate according to digit map", or "notify" action. A single notification can report a list of events that will be reported in the order in which they were detected. The list can only contain persistent events and events that were requested in the **RequestedEvents** parameter of the triggering **NotificationRequest**. Events that were detected on a connection will include the name of that connection. The list will contain the events that were either accumulated (but not notified) or accumulated according to digit map (but no match yet), and the final event that triggered the notification or provided a final match in the digit map. It should be noted that digits are added to the list of observed events as they are accumulated, irrespective of whether they are accumulated according to the digit map or not. For example, if a user enters the digits "1234" and some event E is accumulated between the digits "3" and "4" being entered, the list of observed events would be "1, 2, 3, E, 4".

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

6.3.3 CreateConnection

This command is used to create a connection.

```

ReturnCode
, ConnectionId
[, SpecificEndPointId]
, LocalConnectionDescriptor
[, ResourceID]
    ← CreateConnection(CallId                               , EndpointId
                       [, NotifiedEntity]
                       [, LocalConnectionOptions]
                       , Mode
                       [, RemoteConnectionDescriptor]
                       [, RequestedEvents]
                       [, RequestIdentifier]
                       [, DigitMap]
                       [, SignalRequests]
                       [, QuarantineHandling]
                       [, DetectEvents])

```

This function is used when setting up a connection between two endpoints. A connection is defined by its attributes and the endpoints it associates. The input parameters in CreateConnection provide the data necessary to build one of the two endpoints "view" of a connection.

CallId is a parameter that identifies the call (or session) to which this connection belongs. This parameter is, at a minimum, unique within the collection of Call Agents that control the same gateways; connections that belong to the same call share the same call-id. The call-id can be used to identify calls for reporting and accounting purposes.

EndPointId is the identifier for the endpoint in the gateway where CreateConnection executes. The EndPointId can be specified fully by assigning a non-wildcarded value to the parameter EndPointId in the function call or it can be underspecified by using the "anyone" wildcard convention. If the endpoint is under-specified, the endpoint identifier will be assigned by the gateway and its complete value returned in the Specific EndpointId parameter of the response. In this case, the endpoint assigned MUST be in service and MUST NOT already have any connections on it.

NotifiedEntity is an optional parameter that specifies a new "notified entity" for the endpoint.

LocalConnectionOptions is a structure that describes the characteristics of the media data connection from the point-of-view of the gateway executing CreateConnection. It instructs the endpoint on send and receive characteristics of the media connection. The basic fields contained in LocalConnectionOptions are:

- **Encoding Method:** A list of literal names for the compression algorithm (encoding/decoding method) used to send and receive media on the connection MUST be specified with at least one value. The entries in the list are ordered by preference. The endpoint MUST choose at least one of the codecs, and the codec SHOULD be chosen according to the preference indicated. If the endpoint receives any media on the connection encoded with a different encoding method, it MAY discard it. The endpoint MUST additionally indicate which of the remaining compression algorithms it is willing to support as alternatives - see clause 7.4.1 for details. A list of permissible encoding methods is specified in a separate IPCablecom document.
- **Packetization Period:** The packetization period in milliseconds, as defined in the SDP standard (IETF RFC 2327 [14]), MUST be specified and with exactly one value or a range of values. The value or range only pertains to media sent. A range is specified as two decimal numbers separated by a hyphen. Note that only valid packetization rates within a specified range may be used by the MTA. A list of permissible packetization periods is specified in TS 101 909-3 [6].
- **Echo Cancellation:** Whether echo cancellation should be used on the line side or not (see note 1). The parameter can have the value "on" (when the echo cancellation is requested) or "off" (when it is turned off). The parameter is optional. When the parameter is omitted, the embedded client MUST apply echo cancellation.

NOTE 1: Echo cancellation on the packet side is not supported.

- **Type of Service:** Specifies the class of service that will be used for sending media on the connection by encoding the 8-bit type of service value parameter of the IP header as two hexadecimal digits. The parameter is optional. When the parameter is omitted, a default value of AOH (unless provisioned otherwise) applies corresponding to an IP precedence bits setting of five.

- **Silence Suppression:** Whether silence suppression should be used or not in the send direction. The parameter can have the value "on" (when silence is to be suppressed) or "off" (when silence is not to be suppressed). The parameter is optional. When the parameter is omitted, the default is not to use silence suppression.

The following LocalConnectionOptions fields are used to support Dynamic Quality of Service (D-QoS) - please refer to annex B for further details:

- **D-QoS GateID:** The GateID for the gate that has been setup at the edge router. The Gate-ID is a 32 bit identifier encoded as a string of up to 8 hex characters. This parameter is optional in general, but mandatory when D-QoS resource reservation and/or committal is to be performed. The presence of this parameter implies that D-QoS is to be performed for this command, where as absence implies that D-QoS is not to be performed.
- **D-QoS Resource Reservation:** Allows explicit control over whether D-QoS resource reservation and/or committal should be performed in the send and/or receive direction or not. The parameter is optional and can have one or more of the following values:

Reserve values:

- "SendReserve" Resources are reserved in the send direction only.
- "ReceiveReserve" Resources are reserved in the receive direction only.
- "SendReceiveReserve" Resources are reserved in the send and receive direction.

Commit values:

- "SendCommit" Resources are committed in the send direction only.
- "ReceiveCommit" Resources are committed in the receive direction only.
- "SendReceiveCommit" Resources are committed in the send and receive direction.

The parameter is optional, and multiple values are separated by commas. When D-QoS is to be performed, and the parameter is either omitted or no value is present, resource reservation **MUST** be performed for both the send and receive direction. The resources reserved are determined by the coding parameters applied to the connection, i.e. encoding method, packetization period, silence suppression, ciphersuite, etc. External parameters, such as the use of payload header suppression may affect the amount of resources reserved as well (see TS 101 909-3 [6] for details).

Receive resources can be reserved and committed without having obtained a RemoteConnectionDescriptor, whereas send resources can be reserved, but not committed, until a RemoteConnectionDescriptor is supplied. Note that, as long as a RemoteConnectionDescriptor has not been received, the resources reserved and committed must be based on the codec(s) selected locally. Once a RemoteConnectionDescriptor is received, the list of codec(s) that can actually be used for sending may only contain a subset of these. The list of codecs that can be used for receiving is however unchanged until the endpoint issues a new LocalConnectionDescriptor.

When D-QoS reservation is to be performed, and the parameter is either omitted or no value is present, resources **MUST** by default be committed based on the connection mode as specified in table 2.

Table 2

Connection Mode	D-QoS
"inactive"	Do not commit
"send only", "replicate"	Commit send
"receive only"	Commit receive
"send/receive", "conference", "network loopback", "network continuity test"	Commit send and receive

If a different commit operation is desired, the appropriate commit value is supplied and will be used instead. If a commit operation is to be performed, but no reservation has been made, or an existing reservation does not fully satisfy the resources to be committed (see note 2), a reservation will be made automatically. If a reserve value is specified, but no commit value is specified, a commit operation will not be performed.

NOTE 2: This is not possible for the CreateConnection command but is noted here for completeness. It is possible for the ModifyConnection command however (see clause 6.3.4).

- **ResourceID:** An existing ResourceID for resources already reserved at the edge router. The use of the ResourceID allows separate reservations to reserve the same resource, however only one of the reservations can be active at a given point in time. The ResourceID is a 32-bit identifier encoded as a string of up to 8 hex characters. The parameter is optional.
- **ReserveDestination:** This optional parameter may specify an IPv4 address, optionally followed by a colon and a UDP port number, that is the destination for the resource reservation. When a UDP port number is not specified, a default value of 9 applies. The ReserveDestination is typically used when resource reservation is to be performed, and a RemoteConnectionDescriptor has not yet been provided for the connection. This enables reservations and downstream commits to be sent to the edge router when the source of a media stream is not yet known (see note 3). When a RemoteConnectionDescriptor has been provided, the parameter is ignored.

NOTE 3: Note that this will enable certain theft-of-service scenarios. See the Dynamic Quality of Service Specification (j.dqos) for details.

The following LocalConnectionOptions fields are used to support the IPCablecom security services:

- **RTP ciphersuite** A list of allowable ciphersuites for RTP security in order of preference. The entries in the list are ordered by preference where the first ciphersuite is the preferred choice. In the case that remote SDP is also present in the message, the endpoint MUST first perform an intersection of the RTP ciphersuites in the remote SDP with this list before making a selection (from the resulting intersection). The endpoint MUST choose exactly one of the ciphersuites. The endpoint SHOULD additionally indicate which of the remaining ciphersuites it is willing to support as alternatives (see clause 7.4.1 for details). Each ciphersuite is represented as ASCII strings consisting of two substrings separated by a slash ('/'), where the first substring identifies the authentication algorithm, and the second substring identifies the encryption algorithm. A list of permissible ciphersuites are specified in the IPCablecom Security specification TS 101 909-11 [28].
- **RTCP ciphersuite** A list of ciphersuites for RTCP security in order of preference. The entries in the list are ordered by preference where the first ciphersuite is the preferred choice. In the case that remote SDP is also present in the message, the endpoint MUST first perform an intersection of the RTP ciphersuites in the remote SDP with this list before making a selection (from the resulting intersection). The endpoint MUST choose exactly one of the ciphersuites. The endpoint SHOULD additionally indicate which of the remaining ciphersuites it is willing to support as alternatives. See clause 7.4.1 for details. Each ciphersuite is represented as ASCII strings consisting of two substrings separated by a slash ('/'), where the first substring identifies the authentication algorithm, and the second substring identifies the encryption algorithm. A list of permissible ciphersuites is specified in the IPCablecom Security specification TS 101 909-11 [28].

The embedded client MUST respond with an error (error code 524 - LocalConnectionOptions inconsistency) if any of the above rules are violated. All of the above mentioned default values can be altered by the provisioning process.

RemoteConnectionDescriptor is the connection descriptor for the remote side of a connection, on the other side of the IP network. It includes the same fields as the LocalConnectionDescriptor (not to be confused with LocalConnectionOptions), i.e. the fields that describe a session according to the SDP standard. Clause 7.4, details the supported use of SDP in the NCS profile. This parameter may have a null value when the information for the remote end is not known. This occurs because the entity that builds a connection starts by sending a CreateConnection to one of the two gateways involved. For the first CreateConnection issued, there is no information available about the other side of the connection. This information may be provided later via a ModifyConnection call.

When codecs are changed during a call, small periods of time may exist where the endpoints use different codes. As stated above, embedded clients MAY discard any media received that is encoded with a different codec than what is specified in the LocalConnectionOptions for a connection.

Mode indicates the mode of operation for this side of the connection. The options are "send only", "receive only", "send/receive", "conference", "inactive", "replicate", "network loopback" or "network continuity test". The handling of these modes is specified in the beginning of clause 6.3. Some endpoints may not be capable of supporting all modes. If the command specifies a mode that the endpoint does not support, an error MUST be returned (error code 517 - unsupported mode). Also, if a connection has not yet received a RemoteConnectionDescriptor, an error MUST be returned if the connection is attempted to be placed in any of the modes "send only", "send/receive", "replicate", "conference", "netwloop" or "netwtest" (error code 527 - missing RemoteConnectionDescriptor).

ConnectionId is a parameter returned by the gateway that uniquely identifies the connection within the context of the endpoint in question.

LocalConnectionDescriptor is a parameter returned by the gateway, which is a session description that contains information about, e.g. addresses and RTP ports for "IN" connections as defined in SDP. It is similar to the RemoteConnectionDescriptor, except that it specifies this side of the connection. Clause 7.4, details the supported use of SDP in the NCS profile.

After receiving a "CreateConnection" command that does not include a RemoteConnectionDescriptor parameter, a gateway is in an ambiguous situation for the connection in question. Because it has exported a LocalConnectionDescriptor parameter, it potentially can receive packets on that connection. Because it has not yet received the other gateway's RemoteConnectionDescriptor parameter, it does not know whether the packets it receives have been authorized by the Call Agent. Thus, it must navigate between two risks, i.e. clipping some important announcements or listening to insane data. The behaviour of the gateway is determined by the value of the mode parameter (subject to security):

- If the mode was set to "receive only", the gateway MUST accept the voice signals received on the connection and transmit them through to the endpoint.
- If the mode was set to "inactive", the gateway MUST (as always) discard the voice signals received on the connection.
- Note that when the endpoint does not have a RemoteConnection Descriptor for the connection, the connection can by definition not be in any of the modes "send only", "send/receive", "replicate", "conference", "netwloop" or "netwtest".

The **RequestedEvents**, **RequestIdentifier**, **DigitMap**, **SignalRequests**, **QuarantineHandling**, and **DetectEvents** parameters are all optional. They can be used by the Call Agent to effectively include a notification request that is executed simultaneously with the creation of the connection. If one or more of these parameters is present, the RequestIdentifier MUST be one of them. Thus, the inclusion of a notification request can be recognized by the presence of a RequestIdentifier. The rest of the parameters may or may not be present. If one of the parameters is not present, it MUST be treated as if it was a normal NotificationRequest with the parameter in question being omitted. This may have the effect of cancelling signals and of stop looking for events.

As an example of use, consider a Call Agent that wants to place a call to an embedded client. The Call Agent should:

- ask the embedded client to create a connection, in order to be sure that the user can start speaking as soon as the phone goes off hook;
- ask the embedded client to start ringing;
- ask the embedded client to notify the Call Agent when the phone goes off-hook.

All of the above can be accomplished in a single CreateConnection command by including a notification request with the RequestedEvents parameters for the off-hook event and the SignalRequests parameter for the ringing signal.

When these parameters are present, the creation of the connection and the notification request MUST be synchronized, which means that they are both either accepted or refused. In our example, the CreateConnection must be refused if the gateway does not have sufficient resources or cannot get adequate resources from the local network access. The off-hook notification request must be refused in the glare condition if the user is already off-hook. In this example, the phone must not ring if the connection cannot be established, and the connection must not be established if the user is already off-hook. An error would be returned instead (error code 401 - phone off hook), which informs the Call Agent of the glare condition.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

ResourceID is a D-QoS parameter that may be returned by the gateway. When a successful D-QoS resource reservation is made, the ResourceID provides a handle for the resources reserved.

6.3.4 ModifyConnection

This command is used to modify the characteristics of a gateway's "view" of a connection. This "view" of the call includes both the local connection descriptor, as well as the remote connection descriptor.

```

ReturnCode
[, LocalConnectionDescriptor]
[, ResourceID]
    ← ModifyConnection(CallId
                        , EndpointId
                        , ConnectionId
                        [, NotifiedEntity]
                        [, LocalConnectionOptions]
                        [, Mode]
                        [, RemoteConnectionDescriptor]
                        [, RequestedEvents]
                        [, RequestIdentifier]
                        [, DigitMap]
                        [, SignalRequests]
                        [, QuarantineHandling]
                        [, DetectEvents])

```

The parameters used are the same as in the CreateConnection command, with the addition of a **ConnectionId** that uniquely identifies the connection within the endpoint. This parameter is returned by the CreateConnection command together with the local connection descriptor. It uniquely identifies the connection within the context of the endpoint.

The **EndpointId** MUST be a fully qualified endpoint name. The local name MUST NOT use the wildcard convention.

The ModifyConnection command can be used to affect connection parameters, subject to the same rules and constraints as specified for CreateConnection:

- Provide information on the other end of the connection through the **RemoteConnectionDescriptor**.
- Activate or deactivate the connection by changing the **mode** parameter's value. This can occur at any time during the connection, with arbitrary parameter values. An activation can, for example, be set to the "receive only" mode.
- Change the parameters of the connection through the **LocalConnectionOptions**, for example, by switching to a different coding scheme, changing the packetization period, or modifying the handling of echo cancellation.

The details of D-QoS operation were specified in the CreateConnection command and generally the same rules apply here, except as noted below:

- **D-QoS GateID:** A D-QoS GateID is mandatory when D-QoS operation is required, unless D-QoS operation has previously been done for the connection in question. In the latter case, the previously supplied D-QoS GateID will then be used.
- **D-QoS Resource Reservation:** Allows explicit control over whether D-QoS resource reservation and/or committal should be performed in the send and/or receive direction or not. The parameter is optional and multiple values can be specified. When the parameter is omitted and D-QoS reservation is to be performed, the default is to reserve in both the send and receive direction, unless a suitable reservation for the connection has already been made (see annex B). In that case, a new reservation will not be made. Resources are committed the same way as for CreateConnection, except when changing to "inactive" mode. In that case, the committed resources MUST be lowered to zero. An existing resource reservation is still maintained though.
- **ResourceID:** The parameter is optional. When supplied, it replaces the ResourceID kept by the embedded client for the connection.
- **ReserveDestination:** The parameter is optional. When supplied, it replaces the ReserveDestination kept by the embedded client for the connection. If a RemoteConnectionDescriptor has been supplied for the connection, the parameter is ignored.

The command will only return a **LocalConnectionDescriptor** if the local connection parameters, such as, e.g. RTP ports, etc. are modified. Thus, if, e.g. only the mode of the connection is changed, a LocalConnectionDescriptor will not be returned. If a connection parameter is omitted, e.g. mode or silence suppression, the old value of that parameter will be retained if possible. If a parameter change necessitates a change in one or more *unspecified* parameters, the gateway is free to choose suitable values for the unspecified parameters that must change (see note 1).

NOTE 1: This can for instance happen if a codec change is specified, and the old codec used silence suppression, but the new one does not support it. If, e.g. the packetization period furthermore was not specified, and the new codec supported the old packetization period, the value of this parameter would not change, as a change would not be necessary.

The RTP address information provided in the RemoteConnectionDescriptor specifies the remote RTP address of the receiver of media for the connection. This RTP address information may have been changed by the Call Agent (see note 2). When RTP address information is given to an embedded client for a connection, the embedded client SHOULD only accept media streams (and RTCP) from the RTP address specified as well. Any media streams received from any other addresses SHOULD be discarded. The IPCablecom Security Specification (under development) should be consulted for additional security requirements.

NOTE 2: For instance if media needs to traverse a firewall.

The **RequestedEvents**, **RequestIdentifier**, **DigitMap**, **SignalRequests**, **QuarantineHandling**, and **DetectEvents** parameters are optional. The parameters can be used by the Call Agent to include a notification request that is tied to and executed simultaneously with the connection modification. If one or more of these parameters is supplied, then RequestIdentifier MUST be one of them. For example, when a call is accepted, the calling gateway should be instructed to place the connection in "send/receive" mode and to stop providing ringback tones. This can be accomplished in a single ModifyConnection command by including a notification request with the RequestedEvents parameters for the on-hook event, and an empty SignalRequests parameter, to stop the provision of ringback tones.

When these parameters are present, the connection modification and the notification request MUST be synchronized, which means that they are both either accepted or refused.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

ResourceID is a D-QoS parameter that is returned by the gateway if it performs a resource reservation and obtains a new ResourceID from the edge router. When a successful D-QoS resource reservation is made, the ResourceID provides a handle for the resources reserved.

6.3.5 DeleteConnection (from the Call Agent)

This command is used to terminate a connection. As a side effect, it collects statistics on the execution of the connection.

```
ReturnCode
, Connection-parameters
  ← DeleteConnection(CallId
                    , EndpointId
                    , ConnectionId
                    [, NotifiedEntity]
                    [, RequestedEvents]
                    [, RequestIdentifier]
                    [, DigitMap]
                    [, SignalRequests]
                    [, QuarantineHandling]
                    [, DetectEvents])
```

The endpoint identifier, in this form of the DeleteConnection command, MUST be fully qualified. Wildcard conventions MUST NOT be used.

In the general case where a connection has two ends, this command has to be sent to both gateways involved in the connection. After the connection has been deleted, packet network media streams previously supported by the connection are no longer available. Any media packets received for the old connection are simply discarded and no new media packets for the stream are sent. When one or more D-QoS reservations and/or committals have been made for the connection, the DeleteConnection command will release the resources reserved.

In response to the DeleteConnection command, the gateway returns a list of parameters that describe the status of the connection. These parameters are:

- **Number of packets sent:** The total number of RTP data packets transmitted by the sender since starting transmission on the connection. The count is not reset if the sender changes its synchronization source identifier (SSRC, as defined in RTP) - for example, as a result of a Modify command. The value is zero if, e.g. the connection was always set in "receive only" mode.

- **Number of octets sent:** The total number of payload octets (i.e. not including header or padding) transmitted in RTP data packets by the sender since starting transmission on the connection. The count is not reset if the sender changes its SSRC identifier - for example, as a result of a ModifyConnection command. The value is zero if, e.g. the connection was always set in "receive only" mode.
- **Number of packets received:** The total number of RTP data packets received by the sender since starting reception on the connection. The count includes packets received from different SSRC if the sender used several values. The value is zero if, e.g. the connection was always set in "send only" mode.
- **Number of octets received:** The total number of payload octets (i.e. not including header or padding) transmitted in RTP data packets by the sender since starting transmission on the connection. The count includes packets received from different SSRC if the sender used several values. The value is zero if, e.g. the connection was always set in "send only" mode.
- **Number of packets lost:** The total number of RTP data packets that have been lost since the beginning of reception. This number is defined to be the number of packets expected less the number of packets actually received, where the number of packets received includes any which are late or are duplicates. The count includes packets received from different SSRC if the sender used several values. Thus, packets that arrive late are not counted as lost, and the loss may be negative if there are duplicates. The count includes packets received from different SSRC if the sender used several values. The number of packets expected is defined to be the extended last sequence number received, less the initial sequence number received. The count includes packets received from different SSRC, if the sender used several values. The value is zero if, e.g. the connection was always set in "send only" mode.
- **Interarrival jitter:** An estimate of the statistical variance of the RTP data packet interarrival time measured in ms and expressed as an unsigned integer. The interarrival jitter "J" is defined to be the mean deviation (smoothed absolute value) of the difference "D" in packet spacing at the receiver compared to the sender for a pair of packets. Detailed computation algorithms are found in IETF RFC 1889 [9]. The count includes packets received from different SSRC if the sender used several values. The value is zero if, e.g. the connection was always set in "send only" mode.
- **Average transmission delay:** An estimate of the network latency, expressed in ms. This is the average value of the difference between the NTP timestamp indicated by the senders of the RTCP messages and the NTP timestamp of the receivers, measured when the messages are received. The average is obtained by summing all the estimates and then dividing by the number of RTCP messages that have been received. It should be noted that the correct calculation of this parameter relies on synchronized clocks. Embedded client devices MAY alternatively estimate the average transmission delay by dividing the measured roundtrip time by two.

For a more detailed definition of these variables, please refer to IETF RFC 1889 [9].

The **NotifiedEntity**, **RequestedEvents**, **RequestIdentifier**, **DigitMap**, **SignalRequests**, **QuarantineHandling**, and **DetectEvents** parameters are optional. They can be used by the Call Agent to transmit a notification request that is tied to and executed simultaneously with the deletion of the connection. However, if one or more of these parameters are present, RequestIdentifier MUST be one of them. For example, when a user hangs up the phone, the gateway might be instructed to delete the connection and to start looking for an off-hook event. This can be accomplished in a single DeleteConnection command also by transmitting the RequestedEvents parameter for the off-hook event and an empty SignalRequests parameter.

When these parameters are present, the delete connection and the notification request MUST be synchronized, which means that they are both either accepted or refused.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

6.3.6 DeleteConnection (from the Embedded Client)

In some circumstances, a gateway may have to clear a connection, for example, because it has lost the resource associated with the connection. The gateway can terminate the connection by using a variant of the DeleteConnection command:

```
ReturnCode
  ← DeleteConnection(CallId,
                    EndpointId,
                    ConnectionId,
```

```
Reason-code,
Connection-parameters)
```

The **EndpointId**, in this form of the DeleteConnection command, MUST be fully qualified. Wildcard conventions MUST NOT be used.

The **Reason-code** is a text string starting with a numeric reason-code and optionally followed by a descriptive text string. A list of reason-codes can be found in clause 6.6.

In addition to the **CallId**, **EndpointId**, and **ConnectionId**, the embedded client will also send the connection's parameters, which would have been returned to the Call Agent in response to a DeleteConnection command from the Call Agent. The reason code indicates the cause of the DeleteConnection. When one or more D-QoS reservations and/or committals have been made for the connection, the embedded client will release the resources reserved.

ReturnCode is a parameter returned by the Call Agent. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

6.3.7 DeleteConnection (Multiple Connections From the Call Agent)

A variation of the DeleteConnection function can be used by the Call Agent to delete multiple connections at the same time. The command can be used to delete all connections that relate to a call for an endpoint:

```
ReturnCode
  ← DeleteConnection(CallId,
                    EndpointId)
```

The **EndpointId**, in this form of the DeleteConnection command, MUST NOT use the "any of" wildcard. All connections for the endpoint(s) with the CallId specified will be deleted. The command does not return any individual statistics or call parameters.

DeleteConnection can also be used by the Call Agent to delete all connections that terminate in a given endpoint:

```
ReturnCode
  ← DeleteConnection(EndpointId)
```

In this form of the DeleteConnection command, Call Agents can take advantage of the hierarchical naming structure of endpoints to delete all the connections that belong to a group of endpoints. In this case, part of the "local endpoint name" component of the EndpointId can be specified using the "all" wildcarding convention, as specified in clause 6.1.1. The "any of" wildcarding convention MUST NOT be used. The command does not return any individual statistics or call parameters.

After the connection has been deleted, packet network media streams previously supported by the connection are no longer available. Any media packets received for the old connection are simply discarded and no new media packets for the stream are sent. When one or more D-QoS reservations and/or committals have been made for the connection, the embedded client will release the resources reserved.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

6.3.8 Auditing

The MGCP is based upon a centralized call control architecture where a Call Agent acts as the remote controller of client devices that provide voice interfaces to users and networks. In order to achieve the same or higher levels of availability as the current PSTN, some protocols have implemented mechanisms to periodically "ping" subscribers in order to minimize the time before an individual outage is detected. In this interest, an MGCP-specific auditing mechanism between the embedded clients and the Call Agents in an IPCablecom system is provided to allow the Call Agent to audit endpoint and connection state and to retrieve protocol-specific capabilities of an endpoint.

Two commands for auditing are defined for the embedded clients:

- **AuditEndPoint:** Used by the Call Agent to determine the status of an endpoint.
- **AuditConnection:** Used by the Call Agent to obtain information about a connection.

Network management beyond the capabilities provided by these commands is generally desirable, e.g. information about the status of the embedded client as opposed to individual endpoints. Such capabilities are expected to be supported by the use of the Simple Network Management Protocol (SNMP) and by definition of a MIB for the embedded client, both of which are outside the scope of the present document.

6.3.8.1 AuditEndPoint

The AuditEndPoint command can be used by the Call Agent to find out the status of a given endpoint.

```
{ReturnCode
  [, EndPointIdList]
  [, NumEndPoints]}|
{ReturnCode
  [, RequestedEvents]
  [, DigitMap]
  [, SignalRequests]
  [, RequestIdentifier]
  [, NotifiedEntity]
  [, ConnectionIdentifiers]
  [, DetectEvents]
  [, ObservedEvents]
  [, EventStates]
  [, VersionSupported]
  [, ReasonCode]
  [, Capabilities]}
<-AuditEndPoint(EndPointId
  [, RequestedInfo]|
  {[SpecificEndPointID]
  [, MaxEndPointIDs]}
```

The **EndPointId** identifies the endpoint that is being audited. The "any of" wildcard convention **MUST NOT** be used.

The "all of" wildcard convention can be used to audit a group of endpoints. If this convention is used, the gateway **MUST** return the list of endpoint identifiers that match the wildcard in the **EndPointIdList** parameter, which is simply a list of SpecificEndPointIds - RequestedInfo **MUST NOT** be included in this case. **MaxEndPointIDs** is a numerical value that indicates the maximum number of EndpointIds to return. If additional endpoints exist, the **NumEndPoints** return parameter **MUST** be present and indicate the total number of endpoints that match the EndpointID specified. In order to retrieve the next block of EndpointIDs, the **SpecificEndPointID** is set to the value of the last endpoint returned in the previous EndPointIDList, and the command is issued.

When the wildcard convention is not used, the (possibly empty) **RequestedInfo** describes the information that is requested for the EndpointId specified - the SpecificEndpointID and MaxEndpointID parameters **MUST NOT** be used then. The following endpoint-specific information can then be audited with this command:

- RequestedEvents, DigitMap, SignalRequests, RequestIdentifier, NotifiedEntity, Connection Identifiers, DetectEvents, ObservedEvents, EventStates, VersionsSupported, ReasonCode, MaxMGCPDatagram, and Capabilities.
- If an endpoint is queried about a parameter it does not understand, the endpoint **MUST NOT** generate an error; instead the parameter **MUST** be omitted from the response.

The response will, in turn, include information about each of the items for which auditing information was requested:

- **RequestedEvents:** The current value of RequestedEvents the endpoint is using including the action associated with each event. Persistent events are included in the list.
- **DigitMap:** The digit map the endpoint is using currently.
- **SignalRequests:** A list of the Time-Out signals that are currently active, On/Off signals that are currently "on" for the endpoint (with or without parameter), and any pending Brief signals (see note). Time-Out signals that have timed-out, and currently playing Brief signals are not included. Parameterized signals are reported with the parameters they were applied with.

NOTE: Currently, there should be no pending brief signals.

- **RequestIdentifier:** The RequestIdentifier for the last NotificationRequest received by the endpoint (includes notification request embedded in connection handling primitives). If no notification request has been received, the value zero will be returned.

- **NotifiedEntity:** The current "notified entity" for the endpoint.
- **ConnectionIdentifiers:** A comma-separated list of ConnectionIdentifiers for all connections that currently exist for the specified endpoint.
- **DetectEvents:** The current value of DetectEvents the endpoint is using. Persistent events are included in the list.
- **ObservedEvents:** The current list of observed events for the endpoint.
- **EventStates:** For events that have auditable states associated with them, the event corresponding to the state the endpoint is in, e.g. off-hook in the example line package if the endpoint is off-hook. The definition of the individual events will state if the event in question has an auditable state associated with it.
- **VersionSupported:** A list of protocol versions supported by the endpoint.
- **ReasonCode:** The value of the ReasonCode parameter in the last RestartInProgress or DeleteConnection command issued by the gateway for the endpoint, or the special value 000 if the endpoint's state is nominal.
- **Capabilities:** The capabilities for the endpoint similar to the LocalConnectionOptions parameter and including event packages and connection modes. If there is a need to specify that some parameters, such as e.g. silence suppression, are only compatible with some codecs, then the gateway will return several capability sets.
- **Compression Algorithm:** A list of supported codecs. The rest of the parameters will apply to all codecs specified in this list.
- **Packetization Period:** A single value or a range may be specified.
- **Bandwidth:** A single value or a range corresponding to the range for packetization periods may be specified (assuming no silence suppression).
- **Echo Cancellation:** Whether echo cancellation is supported or not.
- **Silence Suppression:** Whether silence suppression is supported or not.
- **Type of Service:** Whether type of service is supported or not.
- **Event Packages:** A list of event packages supported. The first event package in the list will be the default package.
- **Modes:** A list of supported connection modes.
- **Dynamic Quality of Service:** Whether Dynamic Quality of Service is supported or not.
- **Security:** Whether IPCablecom Security services are supported or not. If supported, the following parameters may be present as well.
- **RTP Ciphersuites:** A list of authentication and encryption algorithms supported for RTP.
- **RTCP Ciphersuites:** A list of authentication and encryption algorithms supported for RTCP.

The Call Agent may then decide to use the AuditConnection command to obtain further information about the connections.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

If no info was requested and the EndpointId refers to a valid fully-specified EndpointId, the gateway simply returns a successful response (return code 200 - transaction executed normally).

It should be noted, that all of the information returned is merely a snapshot. New commands received, local activity, etc. may alter most of the above. For example the hook-state may change before the Call Agent receives the above information.

6.3.8.2 AuditConnection

Auditing of individual connections on an endpoint can be achieved using the AuditConnection command.

```

ReturnCode
[, CallId]
[, NotifiedEntity]
[, LocalConnectionOptions]
[, Mode]
[, RemoteConnectionDescriptor]
[, LocalConnectionDescriptor]
[, ConnectionParameters]
  ← AuditConnection(EndpointId
                    , ConnectionId
                    [, RequestedInfo])

```

The **EndpointId** identifies the endpoint that is being audited-wildcards **MUST NOT** be used. The (possibly empty) **RequestedInfo** describes the information that is requested for the **ConnectionId** within the EndpointId specified. The following connection info can be audited with this command:

```

CallId, NotifiedEntity, LocalConnectionOptions, Mode, ConnectionParameters,
RemoteConnectionDescriptor, LocalConnectionDescriptor.

```

The response will, in turn, include information about each of the items for which auditing info was requested:

- **CallId:** The CallId for the call to which the connection belongs.
- **NotifiedEntity:** The current "notified entity" for the endpoint.
- **LocalConnectionOptions:** The LocalConnectionOptions supplied for the connection.
- **Mode:** The current connection mode.
- **ConnectionParameters:** Current connection parameters for the connection.
- **LocalConnectionDescriptor:** The LocalConnectionDescriptor that the gateway supplied for the connection.
- **RemoteConnectionDescriptor:** The RemoteConnectionDescriptor that was supplied to the gateway for the connection.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

If no information was requested, and the EndpointId refers to a valid endpoint, the gateway simply checks that the connection specified exists and, if so, returns a positive response (return code 200 - transaction executed).

6.3.9 Restart in Progress

The RestartInProgress command is used by the gateway to signal that an endpoint, or a group of endpoints, is taken out of service or is being placed back in service.

```

ReturnCode
[, NotifiedEntity]
[, VersionSupported]
  ← RestartInProgress(EndPointId
                    , RestartMethod
                    [, RestartDelay]
                    [, ReasonCode])

```

The **EndpointId** identifies the endpoints that are taken in or out of service. The "all of" wildcard convention can be used to apply the command to a group of endpoints, for example, all endpoints that are attached to a specified interface, or even all endpoints that are attached to a given gateway. The "any of" wildcard convention **MUST NOT** be used.

The RestartMethod parameter specifies the type of restart:

- A "graceful" restart method indicates that the specified endpoint(s) will be taken out of service after the specified "restart delay". The established connections are not yet affected, but the Call Agent should refrain from establishing new connections, and should try to gracefully tear down any existing connections.

- A "cancel-graceful" restart method indicates that a gateway is cancelling a previously issued "graceful" restart method for the same endpoints. When this command is sent, the gateway will immediately begin to allow the establishment of new connections on these endpoints
- A "forced" restart method indicates that the specified endpoints are taken out of service abruptly. The established connections, if any, are lost.
- A "restart" method indicates that service will be restored on the endpoints after the specified "restart delay". There are no connections that are currently established on the endpoints.
- A "disconnected" method indicates that the endpoint has become disconnected and is now trying to establish connectivity. The "restart delay" specifies the number of seconds the endpoint has been disconnected. Established connections are not affected.

The optional "restart delay" parameter is expressed as a number of seconds. If the number is absent, the delay value should be considered null. In the case of the "graceful" method, a null delay indicates that the Call Agent should simply wait for the natural termination of the existing connections, without establishing new connections. The restart delay is always considered null in the case of the "forced" and "cancel-graceful" methods. A restart delay of null for the "restart" method indicates that service has already been restored. This typically will occur after gateway startup/reboot. To mitigate the effects of a client IP address change, the Call Agent MAY wish to resolve the embedded client's domain name by querying the DNS regardless of the TTL of a current resource record for the restarted embedded client.

Embedded clients SHOULD send a "graceful" or "forced" RestartInProgress message as a courtesy to the Call Agent when they are taken out of service, e.g. by being shutdown, or taken out of service by a network management system, although the Call Agent cannot rely on always receiving such messages. Embedded clients MUST send a "restart" RestartInProgress message with a null delay to their Call Agent when they are back in service according to the restart procedure specified in clause 6.4.3.5 - Call Agents can rely on receiving this message. Also, embedded clients MUST send a "disconnected" RestartInProgress message to their current "notified entity" according to the "disconnected" procedure specified in clause 6.4.3.6. The "restart delay" parameter MUST NOT be used with the "forced" restart method.

The optional ReasonCode parameter may be used to indicate the cause of the restart.

The RestartInProgress message will be sent to the current "notified entity" for the EndpointId in question. It is expected that a default Call Agent, i.e. "notified entity", has been provisioned for each endpoint so, after a reboot, the default Call Agent will be the "notified entity" for each endpoint. Embedded clients MUST take full advantage of wild-carding to minimize the number of RestartInProgress messages generated when multiple endpoints in a gateway restart and the endpoints are managed by the same Call Agent.

ReturnCode is a parameter returned by the Call Agent. It indicates the outcome of the command and consists of an integer number (see clause 6.5) optionally followed by commentary.

A NotifiedEntity MAY additionally be returned with the response to the RestartInProgress from the Call Agent - this should normally only be done in response to "restart" or "disconnected" (see also clauses 6.4.3.5 and 6.4.3.6):

- If the response indicated success (return code 200 - transaction executed), the restart in question completed successfully, and the NotifiedEntity returned is the new "notified entity" for the endpoint(s).
- If the response from the Call Agent indicated an error code, the restart in question is not yet complete. If the response was 521 (endpoint redirected), then the response MUST include a Notified Entity parameter, which specifies the new "notified entity" for the endpoint(s) and MUST be used when retrying the restart in question (as a new transaction).

In the case of RestartMethod parameters of the types "restart" and "disconnected", the restart in question MUST be retried whenever the Call Agent returns a transient (4xx) error code, whereas it SHOULD be retried for any other RestartMethod. It is RECOMMENDED that any type of restart be terminated if a permanent (5xx) error code is returned, except for the case of return code 521 as specified above.

Finally, a **VersionSupported** parameter with a list of supported versions may be returned if the response indicated version incompatibility (error code 528).

6.4 States, failover and race conditions

In order to implement proper call signalling, the Call Agent must keep track of the state of the endpoint, and the gateway must make sure that events are properly notified to the Call Agent. Special conditions may exist when the gateway or the Call Agent are restarted: the gateway may need to be redirected to a new Call Agent during "failover" procedures; similarly, the Call Agent may need to take special action when the gateway is taken offline, or restarted.

6.4.1 Recaps and highlights

As mentioned in clause 6.1.4, Call Agents are identified by their domain name, and each endpoint has one, and only one, "notified entity" associated with it at any given point in time. In this clause we recap and highlight the areas that are of special importance to reliability and fail-over in MGCP:

- A Call Agent is identified by its domain name, not its network addresses, and several network addresses can be associated with a domain name.
- An endpoint has one, and only one, Call Agent associated with it at any given point in time. The Call Agent associated with an endpoint is the current value of the "notified entity".
- The "notified entity" is initially set to a provisioned value. When commands with a NotifiedEntity parameter is received for the endpoint, including wild-carded endpoint-names, the "notified entity" is set to the value specified. If the "notified entity" for an endpoint is empty or has not been set explicitly (see note), the "notified entity" defaults to the source address of the last connection handling command or notification request received for the endpoint. In this case, the Call Agent will thus be identified by its network address, which SHOULD only be done on exceptional basis.

NOTE: This could for instance happen by specifying an empty NotifiedEntity parameter.

- Responses to commands are always sent to the source address of the command, regardless of the current "notified entity". When a Notify message needs to be piggy-backed with the response, the datagram is still sent to the source address of the new command received, regardless of the NotifiedEntity for any of the commands.
- When the "notified entity" refers to a domain name that resolves to multiple IP-addresses, endpoints are capable of switching between each of these addresses, however they cannot change the "notified entity" to another domain name on their own. A Call Agent can however instruct them to switch by providing them with a new "notified entity".
- If a Call Agent becomes unavailable, the endpoints managed by that Call Agent will eventually become "disconnected". The only way for these endpoints to become connected again is either for the failed Call Agent to become available again, or for another (backup) Call Agent to contact the affected endpoints with a new "notified entity".
- When another (backup) Call Agent has taken over control of a group of endpoints, it is assumed that the failed Call Agent will communicate and synchronize with the backup Call Agent in order to transfer control of the affected endpoints back to the original Call Agent, if so desired. Alternatively, the failed Call Agent could simply become the backup Call Agent now.

We should note that handover conflict resolution between separate Call Agents is not provided - we are relying strictly on the Call Agents knowing what they are doing and communicating with each other (although AuditEndpoint can be used to learn about the current "notified entity").

6.4.2 Retransmission and detection of lost associations

The MGCP protocol is organized as a set of transactions, each of which is composed of a command and a response. The MGCP messages, being carried over UDP, may be subject to losses. In the absence of a timely response (see clause 7.5), commands are repeated. Gateways MUST keep in memory a list of the responses that they sent to recent transactions, and a list of the transactions that are currently being executed. Recent is here defined by the value $T_{t_{hist}}$ that specifies the number of seconds that responses to old transactions must be kept for. The default value for $T_{t_{hist}}$ is 30 s.

The transaction identifiers of incoming commands are first compared to the transaction identifiers of the recent responses. If a match is found, the gateway does not execute the transaction, but simply repeats the old response. If a match to a previously responded to transaction is not found, the transaction identifier of the incoming command is compared to the list of transactions that have not yet finished executing. If a match is found, the gateway does not execute the transaction, which is simply ignored - a response will be provided when the execution of the command is complete.

This repetition mechanism is used to guard against four types of possible errors:

- transmission errors, when, e.g. a packet is lost due to noise on a line or congestion in a queue;
- component failure, when, e.g. an interface for a Call Agent becomes unavailable;
- Call Agent failure, when, e.g. all interfaces for a Call Agent becomes unavailable;
- failover, when a new Call Agent is "taking over" transparently.

The elements should be able to derive from the past history an estimate of the packet loss rate. In a properly configured system, this loss rate should be very low, typically less than 1 % on average. If a Call Agent or a gateway has to repeat a message more than a few times, it is very legitimate to assume that something else than a transmission error is occurring. For example, given a uniformly distributed loss rate of 1 %, the probability that 5 consecutive transmission attempts fail is 1 in 100 billion, an event that should occur less than once every 10 days for a Call Agent that processes 1 000 transactions per second. (Indeed, the number of repetitions that is considered excessive should be a function of the prevailing packet loss rate.) When errors are non-uniformly distributed, the consecutive failure probability can become somewhat higher. We should note that the "suspicion threshold", which we will call "Max1", is normally lower than the "disconnection threshold", which we will call "Max2", and which should be set to a larger value.

A classic retransmission algorithm would simply count the number of successive repetitions, and conclude that the association is broken after re-transmitting the packet an excessive number of times (typically between 7 and 11 times). In order to account for the possibility of an undetected or in-progress "failover", we modify the classic algorithm as follows:

- The gateway **MUST** always check for the presence of a new Call Agent. It can be noticed by:
 - receiving a command where the NotifiedEntity points to a new Call Agent; or
 - receiving a redirection response pointing to a new Call Agent.
- If a new Call Agent is detected, the gateway **MUST** direct retransmissions of any outstanding commands for the endpoint(s) redirected to that new Call Agent. Responses to new or old commands are still transmitted to the source address of the command.
- Prior to any retransmission, it is checked that the time elapsed since the sending of the initial datagram is no greater than $T_{s_{max}}$. If more than $T_{s_{max}}$ time has elapsed, the endpoint becomes disconnected.
- If the number of retransmissions to this Call Agent equals "Max1", the gateway **MAY** actively query the name server in order to detect the possible change of Call Agent interfaces, regardless of the Time To Live (TTL) associated with the DNS record.

The gateway may have learned several IP addresses for the Call Agent. If the number of retransmissions for this IP address is greater than or equal to "Max1" and lower than "Max2", and there are more IP addresses that have not been tried, then the gateway **MUST** direct the retransmissions to the remaining alternate addresses in its local list.

- If there are no more interfaces to try, and the number of retransmissions is Max2, then the gateway **SHOULD** contact the DNS one more time to see if any other interfaces have become available. If not, the endpoint(s) managed by this Call Agent are now disconnected. When an endpoint becomes disconnected, it **MUST** then initiate the "disconnected" procedure as specified in clause 6.4.3.6.

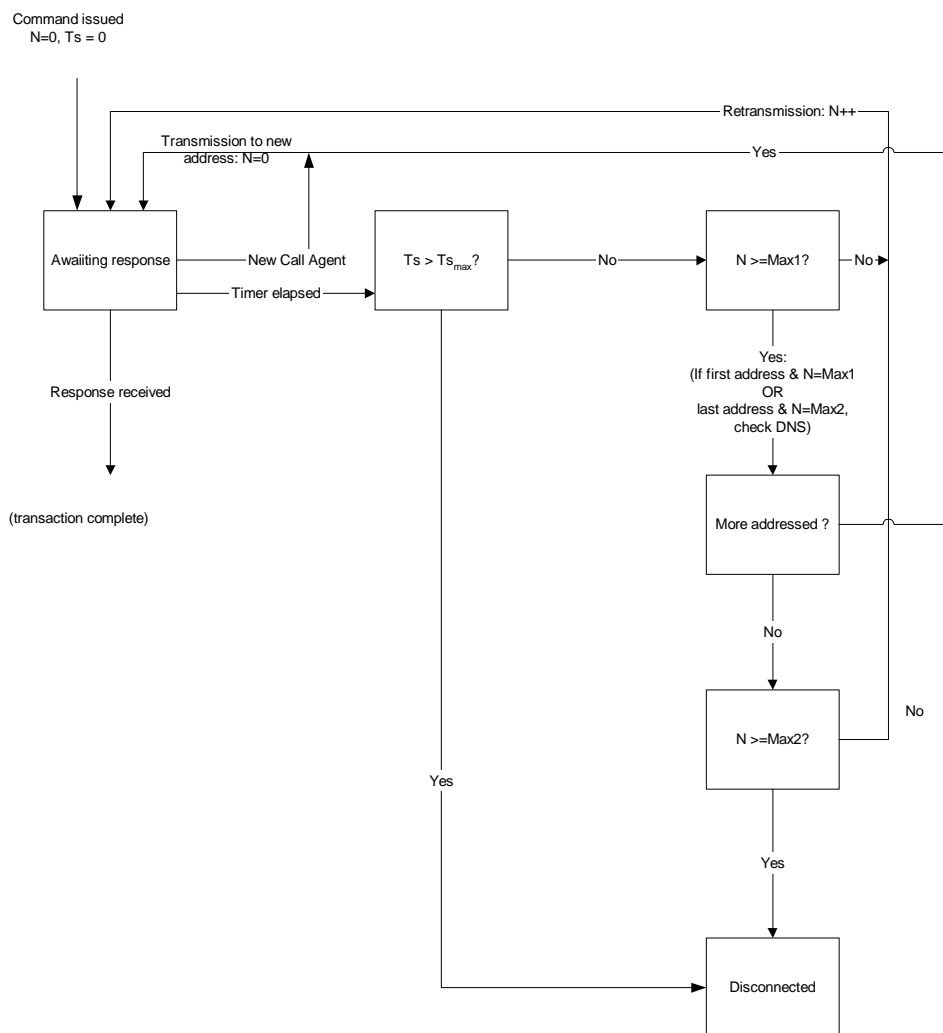


Figure 3

In order to adapt to network load automatically, MGCP specifies exponentially increasing timers (see clause 7.5.2). If the initial time-out is set to 200 ms, the loss of a fifth retransmission will be detected after about 6 s. This is probably an acceptable waiting delay to detect a failover. The retransmissions should continue after that delay not only in order to perhaps overcome a transient connectivity problem, but also in order to allow some more time for the execution of a failover - waiting a total delay of 30 s is probably acceptable.

It should be noted, that there is an intimate relationship between $T_{s_{max}}$, $T_{t_{hist}}$, and the maximum transit time, $T_{p_{max}}$. Specifically, the following relation MUST be satisfied to prevent retransmitted commands from being executed more than once:

$$T_{t_{hist}} \geq T_{s_{max}} + T_{p_{max}}$$

The default value for $T_{s_{max}}$ is 20 s. Thus, if the assumed maximum propagation delay is 10 s, then responses to old transactions must be kept for a period of at least 30 s. The importance of having the sender and receiver agree on these values cannot be overstated.

The default value for Max1 is 5 retransmissions and the default value for Max2 is 7 retransmissions. Both of these values may be altered by the provisioning process.

Furthermore, the provisioning process MUST be able to disable one or both of the Max1 and Max2 DNS queries.

6.4.3 Race conditions

In this clause we describe how MGCP deals with race conditions.

First of all, MGCP deals with race conditions through the notion of a "quarantine list" that quarantines events and through explicit detection of desynchronization, e.g. for mismatched hook-state due to glare for an endpoint.

Secondly, MGCP does not assume that the transport mechanism will maintain the order of commands and responses. This may cause race conditions that may be obviated through a proper behaviour of the Call Agent by a proper ordering of commands.

Finally, in some cases, many gateways may decide to restart operation at the same time. This may occur, for example, if an area loses power or transmission capability during an earthquake or an ice storm. When power and transmission capability are re-established, many gateways may decide to send RestartInProgress commands simultaneously, which could lead to very unstable operation if not carefully controlled.

6.4.3.1 Quarantine list

MGCP controlled gateways will receive notification requests that ask them to watch for a list of events. The protocol elements that determine the handling of these events are the "Requested Events" list, the "Digit Map", the "Quarantine Handling", and the "Detect Events" list.

When the endpoint is initialized, the requested events list only consists of persistent events for the endpoint, and the digit map is assumed empty. After reception of a NotificationRequest command, the gateway starts observing the endpoint for occurrences of the events mentioned in the list, including persistent events.

The events are examined as they occur. The action that follows is determined by the "action" parameter associated to the event in the list of requested events, and also by the digit map. The events that are defined as "accumulate" or "accumulate according to digit map" are accumulated in a list of observed events. The events that are marked as "accumulate according to the digit map" will additionally be accumulated in the "current dial string". This will go on until one event is encountered that triggers a Notify command which will be sent to the current "notified entity".

The gateway, at this point, will transmit the Notify command and will place the endpoint in a "notification state". As long as the endpoint is in this "notification state", the events that are detected on the endpoint are stored in a "quarantine" buffer for later processing. The events are, in a sense, "quarantined". The events detected are the events specified by the union of the RequestedEvents parameter and the most recently received DetectEvents parameter or, in case no DetectEvents parameter has been received, the events that are referred to in the RequestedEvents. Persistent events are detected as well.

The endpoint exits the "notification state" when the response (whether success or failure) to the Notify command is received. The Notify command may be retransmitted in the "notification state", as specified in clause 6.4.2.

If the endpoint is or becomes disconnected (see clause 6.4.2) during this, a response to the Notify command will never be received. The Notify command is then lost and hence no longer considered pending, yet the endpoint is still in the "notification state". Should that occur, completion of the disconnected procedure specified in clause 6.4.3.6 shall then lead the endpoint to exit the "notification state". When the endpoint exits the "notification state" it resets the list of observed events and the "current dial string" of the endpoint to a null value.

Following that point, the behaviour of the gateway depends on the value of the QuarantineHandling parameter in the triggering NotificationRequest command:

If the Call Agent had specified that it expected at most one notification in response to the notification request command ("lockstep" mode), then the gateway **MUST** simply keep on accumulating events in the quarantine buffer until it receives the next notification request command. Until this happens, the endpoint is in a "lockstep state", and events that occur and are to be detected are simply stored in the quarantine buffer. The events to be quarantined are the same as in the "notification state". Once the new NotificationRequest is received and executed successfully, the endpoint exits the "lockstep state".

If, however, the gateway is authorized to send multiple successive Notify commands ("loop" mode), it will proceed as follows. When the gateway exits the "notification state", it resets the list of observed events and the "current dial string" of the endpoint to a null value and starts processing the list of quarantined events, using the already received list of requested events and digit map. When processing these events, the gateway may encounter an event, which triggers a Notify command to be sent. If that is the case, the gateway can adopt one of the two following behaviours.

It can immediately transmit a Notify command that will report all events that were accumulated in the list of observed events until the triggering event, included, leaving the unprocessed events in the quarantine buffer. Or it can attempt to empty the quarantine buffer and transmit a single Notify command reporting several sets of events. The "current dial string" MUST then be reset to a null value after each triggering event. The events that follow the last triggering event MUST be left in the quarantine buffer.

If the gateway transmits a Notify command, the endpoint will re-enter and remain in the "notification state" until the acknowledgement is received (as described above). If the gateway does not find a quarantined event that triggers a Notify command, it places the endpoint in a normal state. Events are then processed as they come, in exactly the same way as if a Notification Request command had just been received.

A gateway can receive at any time a new NotificationRequest command for the endpoint, including the case where the endpoint is disconnected, which will also have the effect of taking the endpoint out of the "notification state" assuming the NotificationRequest executes successfully. Activating an embedded Notification Request is here viewed as receiving a new Notification Request as well.

When a new NotificationRequest is received in the "notification state", the gateway SHOULD attempt to deliver the pending Notify (note that a Notify that was lost due to being disconnected, is no longer considered pending) prior to a successful response to the new NotificationRequest. It does so by using the "piggy-backing" functionality of the protocol and placing the messages (commands and responses) to be sent in order with the oldest message first. The messages will then be sent in a single packet to the source of the new NotificationRequest, regardless of the source and "notified entity" for the old and new command. The steps involved are the following:

- 1) the gateway builds a message that carries in a single packet a repetition of the old outstanding Notify command and the response to the new NotificationRequest command.
- 2) the endpoint is then taken out of the "notification state" without waiting for the response to the Notify command.
- 3) a copy of the outstanding Notify command is kept until a response is received. If a time-out occurs, the Notify will be repeated, in a packet that will also carry a repetition of the response to the NotificationRequest:
 - If the packet carrying the response to the NotificationRequest is lost, the Call Agent will retransmit the NotificationRequest. The gateway will reply to this repetition by retransmitting in a single packet the outstanding Notify command and the response to the NotificationRequest - this datagram will be sent to the source of the NotificationRequest.
 - Notify's for a given endpoint MUST be delivered in-order. If the gateway has to transmit a new Notify before a response to the previous Notify is received, it constructs a packet that piggy-backs a repetition of the old Notify, a repetition of the response to the last NotificationRequest, and the new Notify - this datagram will be sent to current "notified entity". After receiving a NotificationRequest command, the "requested events" list, and "digit map" (if a new one was provided) are replaced by the newly received parameters, and the "current dial string" is reset to a null value. Furthermore, when the NotificationRequest was received in the "notification state", the list of observed events is reset to a null value. The subsequent behaviour is then conditioned by the value of the QuarantineHandling parameter. The parameter may specify that quarantined events, and observed events (which in this case is an empty list), are to be discarded, in which case all quarantined and observed events are discarded. If the parameter specifies that the quarantined and observed events should be processed, the gateway will start processing the list of quarantined and observed events, using the newly received list of "requested events" and "digit map" if provided. When processing these events, the gateway may encounter an event, which triggers a Notify command to be sent. If that is the case, the gateway will immediately transmit a Notify command that will report all events that were accumulated in the list of "observed events" up until and including the triggering event, leaving the unprocessed events in the quarantine buffer. The endpoint then enters the "notification state" again.

A new notification request may be received while the gateway has accumulated events according to the previous notification requests, but has not yet detected any notification-triggering events. The handling of not-yet-notified events is determined, as with the quarantined events, by the quarantine handling parameters:

- If the quarantine-handling parameter specifies that quarantined events shall be ignored, the observed events list is simply reset.
- If the quarantine-handling parameter specifies that quarantined events shall be processed, the observed event list is transferred to the quarantined event list. The observed event list is then reset, and the quarantined event list is processed.

The above procedure applies to all forms of notification requests, regardless of whether they are part of a connection handling command or provided as a NotificationRequest command. Connection handling commands that do not include a notification request are neither affected by nor do they affect the above procedure.

Figure 4 illustrates the procedure specified above assuming all transactions execute successfully.

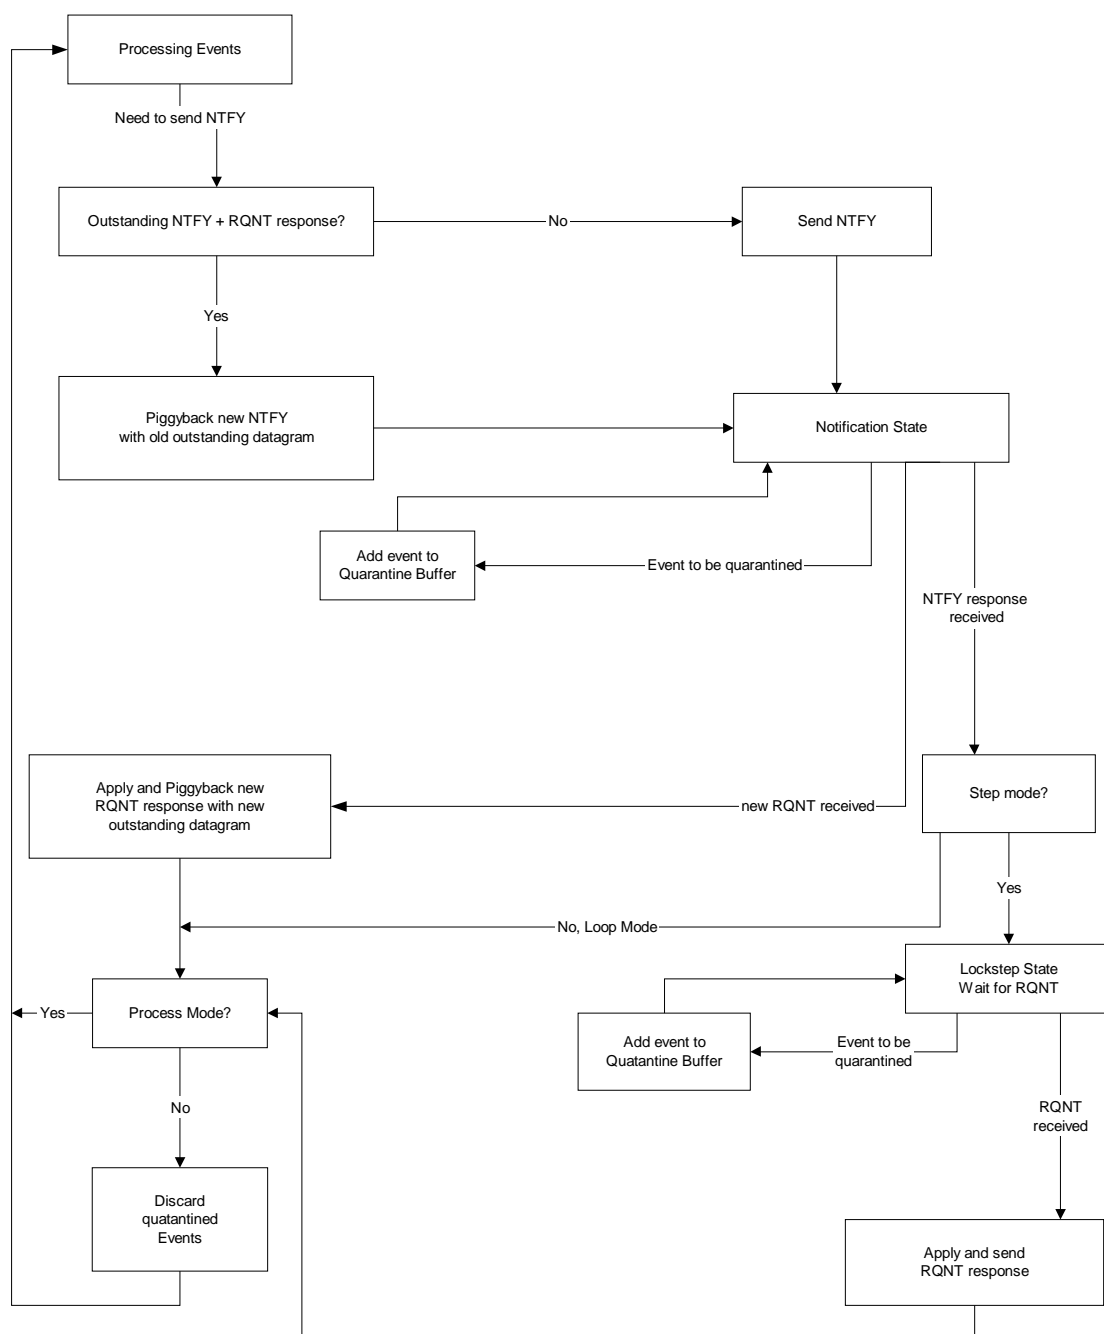


Figure 4

Call Agents SHOULD provide the response to a successful Notify message and the new NotificationRequest in the same datagram using the piggy-backing mechanism.

NOTE: Vendors that choose not to follow the present document should examine Call Agent failure scenarios carefully.

6.4.3.2 Explicit detection

A key element of the state of several endpoints is the position of the hook. Race conditions and state mismatch may occur, for example when the user decides to go off-hook while the Call Agent is in the process of requesting the gateway to look for off-hook events and perhaps apply a ringing signal (the "glare" condition well known in voice-based capabilities).

To avoid this race condition, the gateway MUST check the condition of the endpoint before responding to a NotificationRequest. Specifically, it MUST return an error:

- 1) If the gateway is requested to notify an "off hook" transition while the phone is already off hook (error code 401 - phone off hook).
- 2) If the gateway is requested to notify an "on hook" or "flash hook" condition while the phone is already on hook (error code 402 - phone on hook).

Additionally, individual signal definitions can specify that a signal will only operate under certain conditions, e.g. ringing may only be possible if the phone is already off hook. If such prerequisites exist for a given signal, the gateway MUST return the error specified in the signal definition if the prerequisite is not met.

It should be noted, that the condition check is performed at the time the notification request is received, where as the actual event that caused the current condition may have either been reported, or ignored earlier, or it may currently be quarantined.

The other state variables of the gateway, such as the list of requested events or list of requested signals, are entirely replaced after each successful NotificationRequest, which prevents any long term discrepancy between the Call Agent and the gateway.

When a NotificationRequest is unsuccessful, whether it is included in a connection-handling command or not, the gateway will simply continue as if the command had never been received although an error is returned. As all other transactions, the NotificationRequest MUST operate as an atomic transaction, thus any changes initiated as a result of the command MUST be reverted.

Another race condition can occur when a Notify is issued shortly before the reception by the gateway of a NotificationRequest. The RequestIdentifier is used to correlate Notify commands with NotificationRequest commands thereby enabling the Call Agent to determine if the Notify command was generated before or after the gateway received the new NotificationRequest.

6.4.3.3 Transactional semantics

As the potential transaction completion times increases, e.g. due to external resource reservations, a careful definition of the transactional semantics becomes increasingly important. In particular the issue of race conditions, specifically as it relates to hook-state must be defined carefully.

An important point to consider is, that the hook-state may in fact change between the time a transaction is initiated and the time it completes. More generally, we may say that the successful completion of a transaction depends on one or more pre-conditions where one or more of the pre-conditions may change dynamically during the execution of the transaction.

The simplest semantics for this is simply to require that all pre-conditions MUST be met from the time the transaction is initiated until the transaction completes. Thus, if any of the preconditions change during the execution of the transaction, the transaction MUST fail. Furthermore, as soon as the transaction is initiated, all new events are quarantined. When the outcome of the transaction is known, all quarantined events are then processed.

As an example, consider a transaction that includes a request for the "off-hook" event. When the transaction is initiated the phone is "on-hook" and this pre-condition is therefore met. If the hook-state changes to "off-hook" before the transaction completes, the pre-condition is no longer met, and the transaction therefore immediately fails. The "off-hook" event will now be stored in the "quarantine" buffer which then gets processed.

6.4.3.4 Ordering of commands and treatment of disorder

MGCP does not mandate that the underlying transport protocol guarantees the sequencing of commands sent to a gateway or an endpoint. This property tends to maximize the timeliness of actions, but it has a few drawbacks. For example:

- Notify commands may be delayed and arrive to the Call Agent after the transmission of a new Notification Request command.
- If a new NotificationRequest is transmitted before a response to a previous one is received, there is no guarantee that the previous one will not be received in second position.

Call Agents and gateways that want to guarantee consistent operation of the endpoints can use the rules specified:

- 1) When a gateway handles several endpoints, commands pertaining to the different endpoints can be sent in parallel, for example following a model where each endpoint is controlled by its own process or its own thread.
- 2) When several connections are created on the same endpoint, commands pertaining to different connections can be sent in parallel.
- 3) On a given connection, there should normally be only one outstanding command (create or modify). However, a DeleteConnection command can be issued at any time. In consequence, a gateway may sometimes receive a ModifyConnection command that applies to a previously deleted connection. Such commands MUST be ignored, and an error returned (error code 515 - incorrect connection-id).
- 4) On a given endpoint, there should normally be only one outstanding NotificationRequest command at any time. The RequestId parameter is used to correlate Notify commands with the triggering NotificationRequest.
- 5) In some cases, an implicitly or explicitly wild-carded DeleteConnection command that applies to a group of endpoints can step in front of a pending CreateConnection command. The Call Agent should individually delete all connections whose completion was pending at the time of the global DeleteConnection command. Also, new CreateConnection commands for endpoints named by the wild-carding should not be sent until a response to the wild-carded DeleteConnection command is received.
- 6) When commands are embedded within each other, sequencing requirements for all commands MUST be adhered to. For example a CreateConnection command with a notification request in it must adhere to the sequencing requirements for CreateConnection and NotificationRequest at the same time.
- 7) AuditEndpoint and AuditConnection are not subject to any sequencing.
- 8) RestartInProgress must always be the first command sent by an endpoint as defined by the restart procedure (see clause 6.4.3.5). Any other command or response must be delivered after this RestartInProgress command (piggy-backing allowed).
- 9) When multiple messages are piggy-backed in a single packet, the messages are always processed in order.

Those of the above rules that specify gateway behaviour MUST be adhered to by embedded clients, however the embedded client MUST NOT make any assumptions as to whether Call Agents follow the rules or not. Consequently gateways MUST always respond to commands, regardless of whether they adhere to the above rules or not.

6.4.3.5 Fighting the Restart Avalanche

Let us suppose that a large number of gateways are powered on simultaneously. If they were to all initiate a RestartInProgress transaction, the Call Agent would very likely be swamped, leading to message losses and network congestion during the critical period of service restoration. In order to prevent such avalanches, the following behaviour MUST be followed:

- 1) When a gateway is powered on, it initiates a restart timer to a random value, uniformly distributed between 0 and a provisionable Maximum Waiting Delay (MWD), e.g. 360 s (see below). Care MUST be taken to avoid synchronicity of the random number generation between multiple gateways that would use the same algorithm.

- 2) The gateway then waits for either the end of this timer, the reception of a command from the Call Agent, or the detection of a local user activity, such as for example an off-hook transition on a residential gateway. A pre-existing off-hook condition results in the generation of an off-hook event.
- 3) When the restart timer elapses, when a command is received, or when an activity or pre-existing off-hook condition is detected, the gateway initiates the restart procedure.

The restart procedure simply states that the endpoint **MUST** send a RestartInProgress command to the Call Agent informing it about the restart and furthermore guarantee that the first message (command or response) that the Call Agent sees from this endpoint **MUST** be this RestartInProgress command. The endpoint **MUST** take full advantage of piggy-backing in achieving this. For example, if an off-hook activity occurs prior to the restart timer expiring, a packet containing the RestartInProgress command, and with a piggy-backed Notify command for the off-hook event will be generated. In the case where the restart timer expires without any other activity, the gateway simply sends a RestartInProgress message.

Note that if the RestartInProgress is piggy-backed with the Response (R) to a command received while restarting, then retransmission of the RestartInProgress does not require piggy-backing of the response R. However, while the endpoint is restarting, a resend of the response R does require the RSIP to be piggy-backed to ensure in-order delivery of the two. The restart procedure is complete once a success response has been received. If an error response is received, the subsequent behaviour depends on the error code in question:

- If the error code indicates a transient error (4xx), then the restart procedure **MUST** be initiated again as a new transaction.
- If the error code is 521, the endpoint is redirected, and the restart procedure **MUST** be initiated again as a new transaction. The 521 response includes a NotifiedEntity which is the entity towards which the restart is initiated.
- If the error code is any other permanent error (5xx where xx is not equal to 21), then it is **RECOMMENDED** that the endpoint no longer initiates the restart procedure on its own (until rebooted) unless otherwise specified. If a command is received, the endpoint **MUST** initiate the restart procedure again.

Should the gateway enter the "disconnected" state while carrying out the restart procedure, the disconnected procedure specified in clause 6.4.3.6 **MUST** be carried out, except that a "restart" rather than "disconnected" message is sent during the procedure.

It is expected that each endpoint in a gateway will have a provisionable Call Agent, i.e. "notified entity", to direct the initial restart message towards. When the collection of endpoints in a gateway is managed by more than one Call Agent, the above procedure must be performed for each collection of endpoints managed by a given Call Agent. The gateway **MUST** take full advantage of wild-carding to minimize the number of RestartInProgress messages generated when multiple endpoints in a gateway restart and the endpoints are managed by the same Call Agent.

The value of MWD is a configuration parameter that depends on the type of the gateway. The following reasoning can be used to determine the value of this delay on residential gateways.

Call agents are typically dimensioned to handle the peak hour traffic load, during which, on average, 10 % of the lines will be busy, placing calls whose average duration is typically 3 min. The processing of a call typically involves 5 to 6 transactions between each endpoint and the Call Agent. This simple calculation shows that the Call Agent is expected to handle 5 to 6 transactions for each endpoint, every 30 min on average, or, to put it otherwise, about one transaction per endpoint every 5 to 6 min on average. This suggests that a reasonable value of MWD for a residential gateway would be 10 to 12 min. In the absence of explicit configuration, embedded clients **MUST** use a default value of 600 s for MWD.

6.4.3.6 Disconnected endpoints

In addition to the restart procedure, embedded clients also have a "disconnected" procedure, which is initiated when an endpoint becomes "disconnected" as described in clause 6.4.2. It should here be noted, that endpoints can only become disconnected when they attempt to communicate with the Call Agent. The following steps are followed by an endpoint that becomes "disconnected":

- 1) A "disconnected" timer is initialized to a random value, uniformly distributed between 0 and a provisionable "disconnected" initial waiting delay ($T_{d_{init}}$), e.g. 15 s. Care **MUST** be taken to avoid synchronicity of the random number generation between multiple gateways and endpoints that would use the same algorithm.

- 2) The gateway then waits for either the end of this timer, the reception of a command from the Call Agent, or the detection of a local user activity for the endpoint, such as for example an off-hook transition.
- 3) When the "disconnected" timer elapses, when a command is received, or when a local user activity is detected, the gateway **MUST** initiate the "disconnected" procedure with a new transaction ID for the endpoint. In the case of local user activity, a provisionable "disconnected" minimum waiting delay (Td_{min}) must furthermore have elapsed since the gateway became disconnected or the last time it ended the "disconnected" procedure in order to limit the rate at which the procedure is performed.
- 4) If the "disconnected" procedure still left the endpoint disconnected, the "disconnected" timer is then doubled, subject to a provisionable "disconnected" maximum waiting delay (Td_{max}), e.g. 600 s, and the gateway proceeds with step 2 again.

The "disconnected" procedure is similar to the restart procedure in that it now simply states that the endpoint **MUST** send a RestartInProgress command to the Call Agent informing it that the endpoint was disconnected and furthermore guarantee that the first message (command or response) that the Call Agent now sees from this endpoint **MUST** be this RestartInProgress command. During each initiation of "disconnected" procedure, the command **MUST** observe the normal retransmission and transaction identifiers requirements (see clause 6.4.2). The endpoint **MUST** take full advantage of piggy-backing in achieving this. The Call Agent may then for instance decide to audit the endpoint, or simply clear all connections for the endpoint.

Note that if a disconnected procedure is already in progress when a command is received the existing disconnect procedure **MUST** be terminated and a new procedure **MUST** be started. This is to support a possible call agent redirection.

Also note that if the RestartInProgress is piggy-backed with the Response (R) to a command received while being disconnected, then retransmission of the RestartInProgress does not require piggy-backing of the response R. However, while the endpoint is disconnected, resending the response R does require the RestartInProgress to be piggy-backed as well to ensure the in-order delivery of the two.

The disconnected procedure is complete once a success response has been received. Error responses are handled similarly to the restart procedure (see clause 6.4.3.5). If the "disconnected" procedure is to be initiated again following an error response, the rate-limiting timer considerations specified above still apply.

NOTE: A disconnected endpoint may wish to send a command (besides RestartInProgress) while it is disconnected. Doing so will only succeed once the Call Agent is reachable again, which raises the question of what to do with such a command meanwhile. At one extreme, the endpoint could drop the command right away, however that would not work very well when the Call Agent was in fact available, but the endpoint had not yet completed the 'disconnect procedure. (Consider, for example, the case where a NotificationRequest was just received which immediately resulted in a Notify being generated.) To prevent such scenarios, disconnected endpoint **MUST NOT** blindly drop new commands to be sent for a period of T_{smax} seconds after they receive a non-audit command. One way to satisfy this requirement is to employ a temporary buffering of commands to be sent, however in doing so, the endpoint must ensure that it:

- does not build up a long queue of commands to be sent;
- does not swamp the Call Agent by rapidly sending too many commands once it is connected again.

Buffering commands to T_{smax} seconds and, once the endpoint is connected again, limiting the rate at which buffered commands are sent to one outstanding command per endpoint is considered safe. IF the endpoint is not connected within T_{imex} seconds after the initial send as described in clause 6.4.2.

The present document purposely does not specify any additional behaviour for a disconnected endpoint. Vendors **MAY** for instance choose to provide silence, play reorder tone, or even enable a downloaded wav file to be played on affected endpoints.

The default value for Td_{init} is 15 s, the default value for Td_{min} , is 15 s, and the default value for Td_{max} is 600 s.

6.5 Return codes and error codes

All MGCP commands receive a response. The response carries a return code that indicates the status of the command. The return code is an integer number, for which three value ranges have been defined:

- value 000 indicates a response acknowledgement (see note);
- values between 100 and 199 indicate a provisional response;
- values between 200 and 299 indicate a successful completion;
- values between 400 and 499 indicate a transient error;
- values between 500 and 599 indicate a permanent error.

NOTE: Response acknowledgement is used for provisional responses (see clause 7.8).

The values that have been defined are listed in table 3.

Table 3

Code	Meaning
000	Response acknowledgement.
100	The transaction is currently being executed. An actual completion message will follow later.
200	The requested transaction was executed normally.
250	The connection(s) was deleted.
400	The transaction could not be executed, due to a transient error.
401	The phone is already off hook.
402	The phone is already on hook.
500	The transaction could not be executed because the endpoint is unknown.
501	The transaction could not be executed because the endpoint is not ready.
502	The transaction could not be executed because the endpoint does not have sufficient resources.
508	Unknown or unsupported quarantine handling.
510	The transaction could not be executed because a protocol error was detected.
511	The transaction could not be executed because the command contained an unrecognized extension.
512	The transaction could not be executed because the gateway is not equipped to detect one of the requested events.
513	The transaction could not be executed because the gateway is not equipped to generate one of the requested signals.
514	The transaction could not be executed because the gateway cannot send the specified announcement.
515	The transaction refers to an incorrect connection-id (may have been already deleted).
516	The transaction refers to an unknown call-id.
517	Unsupported or invalid mode.
518	Unsupported or unknown package.
519	Endpoint does not have a digit map.
520	The transaction could not be executed because the endpoint is "restarting".
521	Endpoint redirected to another Call Agent.
522	No such event or signal.
523	Unknown action or illegal combination of actions.
524	Internal inconsistency in LocalConnectionOptions.
525	Unknown extension in LocalConnectionOptions.
526	Insufficient bandwidth.
527	Missing RemoteConnectionDescriptor.
528	Incompatible protocol version.
529	Internal hardware failure.
532	Unsupported value(s) in LocalConnectionOptions.
533	Response too big.
534	Codec negotiation failure

6.6 Reason codes

Reason codes are used by the gateway when deleting a connection to inform the Call Agent about the reason for deleting the connection. They may also be used in a RestartInProgress command, to inform the Call Agent of the reason for the restart. The reason code is an integer number, and the following values have been defined:

Table 4

Code	Meaning
000	Endpoint state is nominal. (This code is used only in response to audit requests.)
900	Endpoint malfunctioning.
901	Endpoint taken out of service.
902	Loss of lower layer connectivity (e.g. downstream sync).
903	QoS resource reservation was lost.

6.7 Use of local connection options and connection descriptors

The normal sequence in setting up a bi-directional connection involves at least 3 steps:

- 1) The Call Agent asks the first gateway to "create a connection" on an endpoint. The gateway allocates resources to that connection, and responds to the command by providing a "session description" (referred to as its LocalConnectionDescriptor). The session description contains the information necessary for another party to send packets towards the newly created connection.
- 2) The Call Agent then asks the second gateway to "create a connection" on an endpoint. The command carries the "Session description" provided by the first gateway (now referred to as the RemoteConnectionDescriptor). The gateway allocates resources to that connection, and responds to the command by providing its own "session description" (LocalConnectionDescriptor).
- 3) The Call Agent uses a "modify connection" command to provide this second "session description" (now referred to as the RemoteConnectionDescriptor) to the first endpoint. Once this is done, communication can proceed in both directions.

When the Call Agent issues a Create or Modify Connection command, there are thus three parameters that determine the media supported by that connection:

- **LocalConnectionOptions:** Supplied by the Call Agent to control the media parameters used by the gateway for the connection. When supplied, the gateway must conform to these media parameters until either the connection is deleted, or a ModifyConnection command is received.
- **RemoteConnectionDescriptor:** Supplied by the Call Agent to convey the media parameters supported by the other side of the connection. When supplied, the gateway must conform to these media parameters until either the connection is deleted, or a Modify Connection command is received.
- **LocalConnectionDescriptor:** Supplied by the gateway to the Call Agent to convey the media parameters it supports for the connection. When supplied, the gateway must honour the media parameters until either the connection is deleted, or the gateway issues a new LocalConnectionDescription.

In determining which codec(s) to provide in the LocalConnectionDescriptor, there are three lists of codecs that a gateway needs to consider:

- A list of codecs provided in the LocalConnectionOptions (either explicitly or implicitly).
- A list of codecs in the RemoteConnectionDescriptor.
- An internal list of codecs that the gateway can support for the connection. A gateway may support one or more codecs for a given connection.

Codec selection (including all relevant media parameters) can then be described by the following steps:

- 1) An approved list of codecs is formed by taking the intersection of the internal list of codecs and codecs allowed by the LocalConnectionOptions. If LocalConnectionOptions were not provided, the approved list of codecs thus contains the internal list of codecs.
- 2) If the approved list of codecs is empty, a codec negotiation failure has occurred and an error response is generated (error code 534 - codec negotiation failure - is recommended.)
- 3) Otherwise, a negotiated list of codecs is formed by taking the intersection of the approved list of codecs and codecs allowed by the RemoteConnectionDescriptor. If a RemoteConnectionDescriptor was not provided, the negotiated list of codecs thus contains the approved list of codecs.
- 4) If the negotiated list of codecs is empty, a codec negotiation failure has occurred and an error response is generated (error code 534 - codec negotiation failure - is recommended).
- 5) Otherwise, codec negotiation has succeeded, and the negotiated list of codecs is returned in the LocalConnectionDescriptor.

Note that both LocalConnectionOptions and the RemoteConnectionDescriptor can contain a list of codecs ordered by preference. When both are supplied, the gateway should adhere to the preferences provided in the LocalConnectionOptions.

7 Media Gateway Control Protocol

The MGCP implements the media gateway control interface as a set of transactions. The transactions are composed of a command and a mandatory response. There are eight types of commands:

- CreateConnection;
- ModifyConnection;
- DeleteConnection;
- NotificationRequest;
- Notify;
- AuditEndpoint;
- AuditConnection;
- RestartInProgress.

The first four commands are sent by the Call Agent to a gateway. The Notify command is sent by the gateway to the Call Agent. The gateway can also send a DeleteConnection as defined in clause 6.3.6. The Call Agent can send either of the Audit commands to the gateway and, finally, the gateway can send a RestartInProgress command to the Call Agent.

7.1 General description

All commands are composed of a Command header, which for some commands may be followed by a session description.

All responses are composed of a Response header, which for some commands may be followed by a session description.

Headers and session descriptions are encoded as a set of text lines, separated by a carriage return and line feed character (or, optionally, a single line-feed character). The headers are separated from the session description by an empty line.

MGCP uses a transaction identifier with a value between 1 and 999999999 to correlate commands and responses. The transaction identifier is encoded as a component of the command header and is repeated as a component of the response header.

7.2 Command header

The command header is composed of:

- a command line identifying the requested action or verb, the transaction identifier, the endpoint towards which the action is requested, and the MGCP protocol version;
- a set of parameter lines composed of a parameter name followed by a parameter value.

Unless otherwise noted or dictated by other referenced standards, each component in the command header is case insensitive. This goes for verbs as well as parameters and values, and all comparisons **MUST** treat upper and lower case as well as combinations of these as being equal.

7.2.1 Command line

The command line is composed of:

- the name of the requested verb;
- the identification of the transaction;
- the name of the endpoint(s) that should execute the command (in notifications or restarts, the name of the endpoint(s) that is issuing the command);
- the protocol version.

These four items are encoded as strings of printable ASCII characters separated by white spaces, i.e. the ASCII space (0x20) or tabulation (0x09) characters. Embedded clients **SHOULD** use exactly one ASCII space separator, however they **MUST** be able to parse messages with additional white space characters.

7.2.1.1 Requested verb coding

Requested verbs are encoded as four letter upper- and/or lower-case ASCII codes (comparisons **MUST** be case insensitive) as defined in table 5.

Table 5

Verb	Code
CreateConnection	CRCX
ModifyConnection	MDCX
DeleteConnection	DLCX
NotificationRequest	RQNT
Notify	NTFY
AuditEndpoint	AUEP
AuditConnection	AUCX
RestartInProgress	RSIP

New verbs may be defined in future versions of the protocol. It may be necessary, for experimental purposes, to use new verbs before they are sanctioned in a published version of this protocol. Experimental verbs should be identified by a four-letter code starting with the letter X (e.g. XPER).

An embedded client that receives a command with an experimental verb it does not support **MUST** return an error (error code 511 - unrecognized extension).

7.2.1.2 Transaction identifiers

Transaction identifiers are used to correlate commands and responses.

An embedded client supports two separate transaction identifier name spaces:

- a transaction identifier name space for sending transactions; and
- a transaction identifier name space for receiving transactions.

At a minimum, transaction identifiers for commands sent to a given embedded client **MUST** be unique for the maximum lifetime of the transactions within the collection of Call Agents that control that embedded client (see clause 7.5). Thus, regardless of the sending Call Agent, embedded clients can always detect duplicate transactions by simply examining the transaction identifier. The coordination of these transaction identifiers between Call Agents is outside the scope of the present document though.

Transaction identifiers for all commands sent from a given embedded client **MUST** be unique for the maximum lifetime of the transactions (see clause 7.5) regardless of which Call Agent the command is sent to. Thus, a Call Agent can always detect a duplicate transaction from an embedded client by the combination of the domain-name of the endpoint and the transaction identifier. The embedded client in turn can always detect a duplicate response acknowledgement by looking at the transaction id(s).

The transaction identifier is encoded as a string of up to nine decimal digits. In the command lines, it immediately follows the coding of the verb.

Transaction identifiers have values between 1 and 999999999. An MGCP entity **MUST NOT** reuse a transaction identifier more quickly than three min after completion of the previous command in which the identifier was used.

7.2.1.3 Endpoint, Call Agent and NotifiedEntity name coding

The endpoint names and Call Agent names are encoded as e-mail addresses, as defined in IETF RFC 821 [7]. In these addresses, the domain name identifies the system where the endpoint is attached, while the left side identifies a specific endpoint on that system. Both components **MUST** be case insensitive.

Examples of such names are:

aaln/1@ncs2.whatever.net	Analogue access line 1 in the embedded client ncs2 in the "Whatever" network.
Call-agent@ca.whatever.net	Call Agent for the "whatever" network.

The name of notified entities is expressed with the same syntax, with the possible addition of a port number, as in:

- `Call-agent@ca.whatever.net:5234`

In case the port number is omitted, the default MGCP port (2427) will be used. Additional detail on endpoint names can be found in clause 6.1.1.

7.2.1.4 Protocol version coding

The protocol version is coded as the keyword "MGCP" followed by a white space and the version number, which again is followed by the profile name "NCS" and a profile version number. The version numbers are composed of a major version number, a dot, and a minor version number. The major and minor version numbers are coded as decimal numbers. The profile version number defined by the present document is 1.0.

The protocol version for the present document **MUST** be encoded as:

- `MGCP 1.0 NCS 1.0`

The "NCS 1.0" portion signals that this is the NCS 1.0 profile of MGCP 1.0.

An entity that receives a command with a protocol version it does not support, **MUST** respond with an error (error code 528 - Incompatible Protocol Version).

7.2.2 Parameter lines

Parameter lines are composed of a parameter name, which in most cases is composed of a single upper-case character, followed by a colon, a white space, and the parameter value. Parameter names and values are still case-insensitive though. The parameters that can be present in commands are defined in table 6.

Table 6

Parameter name	Code	Parameter value
ResponseAck (see note)	K	See description
CallId	C	Hexadecimal string, length MUST NOT exceed 32 characters.
ConnectionId	I	Hexadecimal string, length MUST NOT exceed 32 characters.
NotifiedEntity	N	An identifier, in IETF RFC 821 [7] format, composed of an arbitrary string and of the domain name of the requesting entity, possibly completed by a port number, as in: Call-agent@ca.whatever.net:5234.
RequestIdentifier	X	Hexadecimal string, length MUST NOT exceed 32 characters.
LocalConnectionOptions	L	See description.
Connection Mode	M	See description.
RequestedEvents	R	See description.
SignalRequests	S	See description.
DigitMap	D	A text encoding of a digit map.
ObservedEvents	O	See description.
ConnectionParameters	P	See description.
ReasonCode	E	See description.
SpecificEndPointId	Z	An identifier, in IETF RFC 821 [7] format, composed of an arbitrary string, optionally followed by an "@" followed by the domain name of the embedded client to which this endpoint is attached.
MaxEndPointIds	ZM	Decimal string, length MUST NOT exceed 16 characters.
NumEndPoints	ZN	Decimal string, length MUST NOT exceed 16 characters.
RequestedInfo	F	See description.
QuarantineHandling	Q	See description.
DetectEvents	T	See description.
EventStates	ES	See description.
ResourceID	DQ-RI	See description.
RestartMethod	RM	See description.
RestartDelay	RD	A number of seconds encoded as a decimal number.
Capabilities	A	See description.
VersionSupported	VS	See description.
NOTE: The ResponseAck parameter was not shown in clause 6.3 as transaction identifiers are not visible in our example API. Implementers may choose a different approach.		

The parameters are not necessarily present in all commands. Table 7 provides the association between parameters and commands. The letter M stands for mandatory, O for optional, and F for forbidden.

Table 7

Parameter name	CRCX	MDCX	DLCX	RQNT	NTFY	AUEP	AUCX	RSIP
ResponseAck (see note 2)	O	O	O	O	O	O	O	O
CallId	M	M	O	F	F	F	F	F
ConnectionId	F	M	O	F	F	F	M	F
RequestIdentifier	O	O	O	M	M	F	F	F
LocalConnectionOptions	O	O	F	F	F	F	F	F
Connection Mode	M	O	F	F	F	F	F	F
RequestedEvents	O (see note 1)	O (see note 1)	O (see note 1)	O (see note 1)	F	F	F	F
SignalRequests	O (see note 1)	O (see note 1)	O (see note 1)	O (see note 1)	F	F	F	F
NotifiedEntity	O	O	O	O	O	F	F	F
ReasonCode	F	F	O	F	F	F	F	O
ObservedEvents	F	F	F	F	M	F	F	F
DigitMap	O	O	O	O	F	F	F	F
Connection parameters	F	F	O	F	F	F	F	F
Specific Endpoint Id	F	F	F	F	F	O	F	F
MaxEndPointIds	F	F	F	F	F	O	F	F
NumEndPoints	F	F	F	F	F	F	F	F
RequestedInfo	F	F	F	F	F	O	O	F
QuarantineHandling	O	O	O	O	F	F	F	F
DetectEvents	O	O	O	O	F	F	F	F
EventStates	F	F	F	F	F	F	F	F
ResourceID	F	F	F	F	F	F	F	F
RestartMethod	F	F	F	F	F	F	F	M
RestartDelay	F	F	F	F	F	F	F	O
Capabilities	F	F	F	F	F	F	F	F
VersionSupported	F	F	F	F	F	F	F	F
RemoteConnectionDescriptor	O	O	F	F	F	F	F	F
NOTE 1: The RequestedEvents and SignalRequests parameters are optional in the NotificationRequest. If these parameters are omitted, the corresponding lists will be considered empty. For the connection handling commands, this applies as well when a RequestIdentifier is included.								
NOTE 2: The ResponseAck parameter was not shown in clause 6.3 as transaction identifiers are not visible in our example API. Implementers may choose a different approach.								

Embedded clients and Call Agents SHOULD always provide mandatory parameters before optional ones, however embedded clients MUST NOT fail if the present document is not followed.

If implementers need to experiment with new parameters, for example when developing a new MGCP application, they should identify these parameters by names that begin with the string "X-" or "X+", such as for example:

- X-FlowerOfTheDay: Daisy

Parameter names that start with "X+" are mandatory parameter extensions. A gateway that receives a mandatory parameter extension that it cannot understand MUST respond with an error (error code 511 - unrecognized extension).

Parameter names that start with "X-" are non-critical parameter extensions. A gateway that receives a non-critical parameter extension that it cannot understand can safely ignore that parameter.

It should be noted that experimental verbs are of the form *XABC*, whereas experimental parameters are of the form *X-BC*.

If a parameter line is received with a forbidden parameter, or any other formatting error, the receiving entity should respond with the most specific error code for the error in question. The least specific error code is 510 - protocol error. Commentary text can always be provided.

7.2.2.1 Response acknowledgement

The response acknowledgement parameter (see note) is used to support the three-way handshake described in clause 7.7. It contains a comma separated list of "confirmed transaction-id ranges".

NOTE: The ResponseAck parameter was not shown in clause 6.3 as transaction identifiers are not visible in our example API. Implementers may choose a different approach.

Each "confirmed transaction-id range" is composed of either one decimal number, when the range includes exactly one transaction, or two decimal numbers separated by a single hyphen, describing the lower and higher transaction identifiers included in the range.

An example of a response acknowledgement is:

- K: 6234-6255, 6257, 19030-19044

7.2.2.2 RequestIdentifier

The request identifier correlates a Notify command with the NotificationRequest that triggered it. A RequestIdentifier is a hexadecimal string; length MUST NOT exceed 32 characters. The string "0" is reserved for reporting of persistent events in the case where no NotificationRequest has been received yet (see clause 6.3.2).

7.2.2.3 Local connection options

The local connection options describe the operational parameters that the Call Agents instructs the gateway to use for a connection. These parameters are:

- The packetization period in milliseconds, encoded as the keyword "p" followed by a colon and a decimal number or a range. A range is specified as two decimal numbers separated by a hyphen.
- The literal name of the compression algorithm, encoded as the keyword "a" followed by a colon and a character string.
- The echo-cancellation parameter, encoded as the keyword "e" followed by a colon and the value "on" or "off".
- The type of service parameter, encoded as the keyword "t" followed by a colon and the value encoded as two hexadecimal digits.
- The silence suppression parameter, encoded as the keyword "s" followed by a colon and the value "on" or "off".

The LocalConnectionOptions parameters used for Dynamic Quality of Service are:

- The D-QoS GateID encoded as the keyword "dq-gi" followed by a colon and a string of up to 8 hex characters corresponding to a 32 bit identifier for the GateID.
- The D-QoS Resource Reservation parameter encoded as the keyword "dq-rr" followed by a colon and a character string. A list of values may be specified in which case the values will be separated by a semicolon. The possible values are:

Table 8

Mode	Meaning
sendresv	Reserve in the send direction only.
recvresv	Reserve in the receive direction only.
snrcresv	Reserve in the send and receive direction.
sendcomt	Commit in the send direction only.
recvcomt	Commit in the receive direction only.
snrccomt	Commit in the send and receive direction.

- The ResourceID encoded as the keyword "dq-ri" followed by a colon and a string of up to 8 hex characters corresponding to a 32 bit identifier for the ResourceID.

- The ReserveDestination is encoded as the keyword "dq-rd" followed by a colon and an IP-address encoded similarly to an IP-address for the domain name portion of an endpoint name. The ReserveDestination may optionally be followed by a colon and up to 5 decimal characters for a UDP port number to use.

The LocalConnectionOptions parameters used for Security are encoded as follows:

- The RTP ciphersuite is encoded as the keyword "sc-rtp" followed by a colon and an RTP ciphersuite string as defined below. A list of values may be specified in which case the values will be separated by a semicolon.
- The RTCP ciphersuite is encoded as the keyword "sc-rtcp" followed by a colon and an RTCP ciphersuite string as defined below. A list of values may be specified in which case the values will be separated by a semicolon.

The RTP and RTCP ciphersuite strings follow the grammar:

```

ciphersuite =                [AuthenticationAlgorithm] "/" [EncryptionAlgorithm]
AuthenticationAlgorithm =    1*( ALPHA / DIGIT / "-" / "_" )
EncryptionAlgorithm =        1*( ALPHA / DIGIT | "-" / "_" )

```

where ALPHA, and DIGIT are defined in IETF RFC 2234 [12]. Whitespaces are not allowed within a ciphersuite. The following example illustrates the use of ciphersuite:

```
62/51
```

The actual list of IPCablecom supported ciphersuites to be provided in the IPCablecom Security Specification.

When several parameters are present, the values are separated by a comma. It MUST be considered an error to include a parameter without a value (error code 524 - LocalConnectionOptions inconsistency).

Examples of local connection options are:

```

L: p:10, a:PCMU
L: p:10, a:PCMU, e:off, t:20, s:on
L: p:30, a:G729A, e:on, t:A0, s:off

```

The type of service hex value "20" implies an IP precedence of 1, and a type of service hex value of "A0" implies an IP precedence of 5.

This set of attributes may be extended by extension attributes. Extension attributes are composed of an attribute name, followed by a colon, and a semicolon separated list of attribute values. The attribute name MUST start with the two characters "x+", for a mandatory extension, or "x-", for a non-mandatory extension. If a gateway receives a mandatory extension attribute that it does not recognize, it MUST reject the command with an error (error code 525 - Unknown extension in LocalConnectionOptions).

7.2.2.4 Capabilities

Capabilities inform the Call Agent about its capabilities when audited. The encoding of capabilities is based on the Local Connection Options encoding for the parameters that are common to both. In addition, capabilities can also contain a list of supported packages, and a list of supported modes.

The parameters used are:

- The packetization period in ms, encoded as the keyword "p" followed by a colon and a decimal number. A range may be specified as two decimal numbers separated by a hyphen.
- The literal name of the compression algorithm, encoded as the keyword "a" followed by a colon and a character string. A list of values may be specified in which case the values will be separated by a semicolon.
- The bandwidth in kilobits per second (1 000 bits per second), encoded as the keyword "b" followed by a colon and a decimal number. A range may be specified as two decimal numbers separated by a hyphen.
- The echo-cancellation parameter, encoded as the keyword "e" followed by a colon and the value "on" if echo cancellation is supported, "off" otherwise.
- The type of service parameter, encoded as the keyword "t" followed by a colon and the value "0" if type of service is not supported, all other values indicate support for type of service.

- The silence suppression parameter, encoded as the keyword "s" followed by a colon and the value "on" if silence suppression is supported, "off" otherwise.
- The event packages supported by this endpoint encoded as the keyword "v" followed by a colon and then a semicolon-separated list of package names supported. The first value specified will be the default package for the endpoint.
- The connection modes supported by this endpoint encoded as the keyword "m" followed by a colon and a semicolon-separated list of connection modes supported as defined in clause 7.2.2.7.
- The keyword "dq-gi" if Dynamic Quality of Service is supported.
- The keyword "sc-rtcp" followed by a colon and a semi-colon separated list of RTP ciphersuites, using the same encoding as in the LocalConnectionOptions.
- The keyword "sc-rtcp" followed by a colon and a semi-colon separated list of RTCP ciphersuites, using the same encoding as in the LocalConnectionOptions.

When several parameters are present, the values are separated by a comma.

Examples of capabilities are:

```
A: a:PCMU;G729A, p:10-100, e:on, s:off, v:L;S, m:sendonly;recvonly;sendrecv;inactive
A: a:G729A; p:30-90, e:on, s:on, v:L;S,
m: sendonly;recvonly;sendrecv;inactive;confrnce,
dq-gi, sc-rtcp: 64/51;60/51, sc-rtcp: 2/3;1/3
```

Note that the codecs and security algorithms are merely examples - separate IP-Cablecom specifications detail the actual codecs and algorithms supported, as well as the encoding used (see TS 101 909-3 [6] and TS 101 909-11 [28]). Note also that each set of capabilities is provided on a single line. The examples above show each set on multiple lines due only to formatting restraints of the present document.

7.2.2.5 Connection parameters

Connection parameters are encoded as a string of type and value pairs, where the type is a two-letter identifier of the parameter, and the value a decimal integer. Types are separated from values by an "=" sign. Parameters are separated from each other by a comma.

The connection parameter types are specified in table 9.

Table 9

Connection parameter name	Code	Connection parameter value
Packets Sent	PS	The number of packets that were sent on the connection.
Octets Sent	OS	The number of octets that were sent on the connection.
Packets Received	PR	The number of packets that were received on the connection.
Octets Received	OR	The number of octets that were received on the connection.
Packets Lost	PL	The number of packets that were not received on the connection, as deduced from gaps in the sequence number.
Jitter	JI	The average inter-packet arrival jitter, in ms, expressed as an integer number.
Latency	LA	Average latency, in ms, expressed as an integer number.

Extension connection parameter names are composed of the string "X-" followed by a two letters extension parameter name. Call Agents that receive unrecognized extensions MUST silently ignore these extensions.

An example of a connection parameter encoding is:

- P: PS=1245, OS=62345, PR=0, OR=0, PL=0, JI=0, LA=48

7.2.2.6 Reason codes

Reason codes are three-digit numeric values. The reason code is optionally followed by a white space and commentary, e.g.:

- 900 Endpoint malfunctioning

A list of reason codes can be found in clause 6.6.

7.2.2.7 Connection mode

The connection mode describes the connection's operation mode. The possible values are:

Table 10

Mode	Meaning
M: sendonly	The gateway should only send packets.
M: recvonly	The gateway should only receive packets.
M: sendrecv	The gateway should send and receive packets.
M: confrnce	The gateway should send and receive packets according to conference mode.
M: inactive	The gateway should neither send nor receive packets.
M: replicate	The gateway should only send packets according to replicate mode.
M: netwloop	The gateway should place the endpoint in Network Loopback mode.
M: netwtest	The gateway should place the endpoint in Network Continuity Test mode.

7.2.2.8 Event/signal name coding

Event/signal names are composed of an optional package name, separated by a slash (/) from the name of the actual event. The event name can optionally be followed by an at sign (@) and the identifier of a connection on which the event should be observed. Event names are used in the RequestedEvents, SignalRequests, DetectEvents, ObservedEvents, and EventStates parameters. Each event is identified by an event code. These ASCII encodings are not case sensitive. Values such as "hu", "Hu", "HU" or "hU" should be considered equal.

The following are examples of event names:

X/hu	On-hook transition, in the example line package.
X/0	Digit 0 in the example line package.
hf	Flash-hook, assuming that the example line package is the default package for the endpoint.
X/rt@0A3F58	Ringback on connection "0A3F58".

In addition, the range and wildcard notation of events can be used, instead of individual names, in the RequestedEvents and DetectEvents (but not SignalRequests ObservedEvents, or EventStates):

X/[0-9]	Digits 0 to 9 in the example line package.
X/X	Digits 0 to 9 in the example line package.
[0-9*#A-D]	All digits and letters in the example line package (default for endpoint).
X/all	All events in the example line package.

Finally, the star sign can be used to denote "all connections", and the dollar sign can be used to denote the "current" connection. The following are examples of such notations:

X/rt@*	Ringback on all connections for the endpoint.
X/rt@\$	Ringback on the current connection.

An initial set of event packages for embedded clients can be found in annex A.

7.2.2.9 RequestedEvents

The RequestedEvents parameter provides the list of events that have been requested. The currently defined event codes are described in annex A.

Each event can be qualified by a requested action, or by a list of actions. Not all actions can be combined - please refer to clause 6.3.1 for valid combinations. The actions, when specified, are encoded as a list of keywords enclosed in parenthesis and separated by commas. The codes for the various actions are:

Table 11

Action	Code
Notify immediately	N
Accumulate	A
Accumulate according to digit map	D
Ignore	I
Keep Signal(s) active	K
Embedded NotificationRequest	E
Embedded ModifyConnection	C

If a digit map is not provided when the "accumulate according to digit map" action is specified, the endpoint simply uses its current digit map. If the endpoint does not have any digit maps currently, an error must be returned (error code 519 - no digit map).

When no action is specified, the default action is to notify the event. This means that, for example, "oc" and "oc(N)" are equivalent. Events that are not listed are discarded, except for persistent events.

The digit-map action can only be specified for the digits, letters, and timers.

The requested events list is encoded on a single line, with event/action groups separated by commas. Examples of RequestedEvents encodings are (using the example line package):

```
R: hu(N), hf(N) Notify on-hook, notify hook-flash.
R: hu(N), [0-9#T](D) Notify on-hook, accumulate digits according to digit map.
```

The embedded NotificationRequest follows the format:

```
E ( R( <RequestedEvents> ), D( <Digit Map> ), S( <SignalRequests> ) )
```

with each of R, D, and S being optional and possibly supplied in another order. The following example illustrates the use of Embedded NotificationRequest with the example line package:

```
R: hd(A, E(S(d1), R( B/oc(N), [0-9#T](D) ), D((1xxxxxxxxxxx|9011x.T)) ) )
```

On off-hook, accumulate the event, provide dial-tone and start accumulating digits according to the digit map supplied. Stop dial-tone when the first digit is input, or, if no digit is input before the dial-tone times out. Notify the operation complete. Otherwise, notify the off-hook and collected digits when a match, mismatch, or inter-digit timeout has occurred. It should be noted, that since on-hook is a persistent event, it will still be detected and notified although it has not been specified here.

The embedded ModifyConnection action follows the format:

```
C(M(<ConnectionModel>( <ConnectionID1> ) ) , ... ,
M(<ConnectionModen>(ConnectionIDn ) ) )
```

The following example illustrates the use of Embedded ModifyConnection with the example line package:

```
R: hf(A, C(M(inactive(X43DC)), M(sendrecv($)))) , B/oc(N), B/of(N)
```

On hook-flash, change the connection mode of connection "X43DC" to "inactive", and then change the connection mode of the "current connection" to "send receive". Notify events on "operation complete" and "operation failure".

7.2.2.10 SignalRequests

The SignalRequests parameter provides the name of the signals that have been requested. The currently defined signals can be found in annex A. A given signal can only appear once in the list, and all signals will, by definition, be applied at the same time.

Some signals can be qualified by signal parameters. When a signal is qualified by multiple signal parameters, the signal parameters are separated by commas. Each signal parameter MUST follow the format specified below (white spaces allowed):

```
signal-parameter      =  signal-parameter-value |  signal-parameter-name "="signal-parameter-
                        value | signal-parameter-name "(" signal-parameter-list ")"
signal-parameter-list =  signal-parameter-value 0*( "," signal-parameter-value )
```

where signal-parameter-value may be either a string or a quoted string, i.e. a string surrounded by two double quotes. Two consecutive double-quotes in a quoted string will escape a double-quote within that quoted string. For example, "ab" "c" will produce the string ab" c .

Each signal has one of the following signal-types associated with it (see clause 6.3.1):

- On/Off (OO);
- TimeOut (TO);
- BRief (BR).

On/Off signals can be parameterized with a "+" to turn the signal on, or a "-" to turn the signal off. If an on/off signal is not parameterized, the signal is turned on. Both of the following will turn the vmwi signal from the example line package on:

```
vmwi(+), vmwi.
```

Time-out signals can be parameterized with the signal parameter "TO" and a time-out value that overrides the default time-out value. If a time-out signal is not parameterized with a time-out value the default time-out value will be used. Both of the following will apply the ringing signal from the example line package for 6 s:

```
rg(to=6000);
rg(to(6000)).
```

Individual signals may define additional signal parameters.

The signal parameters will be enclosed within parenthesis as shown above.

When several signals are requested, their codes are simply separated by a comma.

7.2.2.11 ObservedEvents

The observed events parameters provide the list of events that have been observed. The event codes are the same as those used in the NotificationRequest. When an event is detected on a connection, the observed event will identify the connection the event was detected on using the "@<connection>" syntax. Examples of observed events using the example line package are:

```
O: hu
O: 8,2,9,5,5,5,5,T
O: hf,hf,hu
```

Events that have been accumulated according to digit map, are reported as individual events in the order they were detected. Other events may be mixed in between them. It should be noted that if the "current dial string" is non-empty with a partial match, and another event occurs that results in a Notify message being generated, the partially matched "current dial string" will be included in the list of observed events, and the "current dial string" will then be cleared - please refer to clause 6.4.3.1 for details.

7.2.2.12 RequestedInfo

The RequestedInfo parameter contains a comma separated list of parameter codes, as defined in the "Parameter lines" clause - clause 6.3.8 lists the parameters that can be audited. The following values are supported as well:

RequestedInfo Parameter	Code
LocalConnectionDescriptor	LC
RemoteConnectionDescriptor	RC

For example, if one wants to audit the value of the NotifiedEntity, RequestIdentifier, RequestedEvents, SignalRequests, DigitMap, DetectEvents, EventStates, LocalConnectionDescriptor, and RemoteConnectionDescriptor parameters, the value of the RequestedInfo parameter will be:

F: N,X,R,S,D,T,ES,LC,RC

The capabilities request, for the AuditEndPoint command, is encoded by the parameter code "A", as in:

F: A

7.2.2.13 QuarantineHandling

The quarantine handling parameter contains a list of comma separated keywords:

- The keyword "process" or "discard" to indicate the treatment of quarantined and observed events. If neither process nor discard is present, process is assumed.
- The keyword "step" or "loop" to indicate whether at most one notification is expected, or whether multiple notifications are allowed. If neither "step" nor "loop" is present, "step" is assumed. Support for these two keywords is optional.

The following values are valid examples:

- Q:loop
- Q:process
- Q:discard,loop

7.2.2.14 DetectEvents

The DetectEvents parameter is encoded as a comma separated list of events, such as for example:

T: hu,hd,hf,[0-9#*]

It should be noted, that no actions can be associated with the events.

7.2.2.15 EventStates

The EventStates parameter is encoded as a comma separated list of events, such as for example:

ES: hu

It should be noted, that no actions can be associated with the events.

7.2.2.16 ResourceID

The ResourceID parameter is a return parameter used for Dynamic Quality of Service to signal the resource ID assigned for the gate in question. The ResourceID is encoded as a string of up to 8 hex characters, such as for example:

DQ-RI: AB345DC

7.2.2.17 RestartMethod

The RestartMethod parameter is encoded as one of the keywords "graceful", "forced", "restart", or "disconnected", as for example:

```
RM: restart
```

7.2.2.18 VersionSupported

The VersionSupported parameter is encoded as a comma separated list of versions supported, such as for example:

```
VS: MGCP 1.0, MGCP 1.0 NCS 1.0
```

7.3 Response header formats

The response header is composed of a response line optionally followed by headers that encode the response parameters.

The response line starts with the response code, which is a three-digit numeric value. The code is followed by a white space, the transaction identifier, and optional commentary preceded by a white space, e.g.:

```
200 1201 OK
```

Table 12 summarizes the response parameters whose presence is mandatory or optional in a response header, as a function of the command that triggered the response assuming the command succeeded. The reader should still study the individual command definitions though as this table only provides summary information. The letter M stands for mandatory, O for optional and F for forbidden.

Table 12

Parameter name	CRCX	MDCX	DLCX	RQNT	NTFY	AUEP	AUCX	RSIP
ResponseAck (note 2)	O (note 1)	O (note 1)	O (note 1)	O (note 1)	O (note 1)	O (note 1)	O (note 1)	O (note 1)
CallId	F	F	F	F	F	F	O	F
ConnectionId	M	F	F	F	F	O	F	F
RequestIdentifier	F	F	F	F	F	O	F	F
LocalConnectionOptions	F	F	F	F	F	O	O	F
Connection Mode	F	F	F	F	F	F	O	F
RequestedEvents	F	F	F	F	F	O	F	F
SignalRequests	F	F	F	F	F	O	F	F
NotifiedEntity	F	F	F	F	F	O	O	O
ReasonCode	F	F	F	F	F	O	F	F
ObservedEvents	F	F	F	F	F	O	F	F
DigitMap	F	F	F	F	F	O	F	F
ConnectionParameters	F	F	O	F	F	F	O	F
Specific Endpoint ID	O	F	F	F	F	O	F	F
MaxEndPointIds	F	F	F	F	F	F	F	F
NumEndPoints	F	F	F	F	F	O	F	F
RequestedInfo	F	F	F	F	F	F	F	F
QuarantineHandling	F	F	F	F	F	F	F	F
DetectEvents	F	F	F	F	F	O	F	F
EventStates	F	F	F	F	F	O	F	F
ResourceID	O	O	F	F	F	F	F	F
RestartMethod	F	F	F	F	F	F	F	F
RestartDelay	F	F	F	F	F	F	F	F
Capabilities	F	F	F	F	F	O	F	F
VersionSupported	F	F	F	F	F	O	F	O
LocalConnection Descriptor	M	O	F	F	F	F	O	F
RemoteConnection Descriptor	F	F	F	F	F	F	O	F
NOTE 1: The ResponseAck parameter MUST NOT be used with any other responses than a final response issued after a provisional response for the transaction in question. In that case, the presence of the ResponseAck parameter MUST trigger a Response Acknowledgement message - any ResponseAck values provided will be ignored.								
NOTE 2: Each endpoint may be provisioned with a separate Call Agent address and port.								

The response parameters are described for each of the commands in the following.

7.3.1 CreateConnection

In the case of a CreateConnection message, the response line is followed by a Connection-Id parameter with a successful response (code 200). A LocalConnectionDescriptor is furthermore transmitted with a positive response. The LocalConnectionDescriptor is encoded as a "session description", as defined in clause 7.4. It is separated from the response header by an empty line, e.g.:

```

200 1204 OK
I: FDE234C8

v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=audio 3456 RTP/AVP 96
a=rtpmap:96 G726-32/8000

```

When a provisional response has been issued previously, the final response may furthermore contain the Response Acknowledgement parameter, and when Dynamic Quality of Service is used, the final response may also contain a ResourceID, as in:

```

200 1204 OK
K:
I: FDE234C8
DQ-RI: 23DB4A43

v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=audio 3456 RTP/AVP 96
a=rtpmap:96 G726-32/8000

```

The final response is acknowledged by a Response Acknowledgement:

```
000 1204
```

7.3.2 ModifyConnection

In the case of a successful ModifyConnection message, the response line is followed by a LocalConnectionDescriptor, if the modification resulted in a modification of the session parameters (e.g. changing only the mode of a connection does not alter the session parameters). The LocalConnectionDescriptor is encoded as a "session description", as defined in clause 7.4. It is separated from the response header by an empty line.

```

200 1207 OK

v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=audio 3456 RTP/AVP 0

```

The response may also contain a ResourceID when Dynamic Quality of Service is used as in:

```

200 1207 OK
DQ-RI: 12345

```

When a provisional response has been issued previously, the final response may furthermore contain the Response Acknowledgement parameter as in:

```

526 1207 No bandwidth
K:

```

The final response is acknowledged by a Response Acknowledgement:

```
000 1207 OK
```

7.3.3 DeleteConnection

Depending on the variant of the DeleteConnection message, the response line may be followed by a Connection Parameters parameter line, as defined in clause 7.2.2.5.

```

250 1210 OK
P: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48

```

7.3.4 NotificationRequest

A NotificationRequest response does not include any additional response parameters.

7.3.5 Notify

A Notify response does not include any additional response parameters.

7.3.6 AuditEndpoint

In the case of an AuditEndPoint the response line may be followed by information for each of the parameters requested - each parameter will appear on a separate line. Parameters for which no value currently exists, e.g. digit map, will still be provided.

Each local endpoint name "expanded" by a wildcard character will appear on a separate line using the "SpecificEndPointId" parameter code, e.g.:

```
200 1200 OK
Z: aaln/1@rgw.whatever.net
Z: aaln/2@rgw.whatever.net
```

An example of a response to an AuditEndPoint message containing a non-wildcarded endpoint name is shown below. Note that the SpecificEndPointId is not provided in this case. Note also that each set of capabilities is provided on a single line. The example below shows each set on multiple lines due only to the formatting restraints of the present document.

```
200 1200 OK
A: a:PCMU;G728, p:10-100, e:on, s:off, t:l, v:X;B, m:sendonly;recvonly;sendrecv;inactive
A: a:G729A; p:30-90, e:on, s:on, t:l, v:X;B, m:sendonly;recvonly;sendrecv;inactive;confrnce
```

7.3.7 AuditConnection

In the case of an AuditConnection, the response may be followed by information for each of the parameters requested. Parameters for which no value currently exists will still be provided. Connection descriptors will always appear last and each will be preceded by an empty line, as for example:

```
200 1203 OK
C: A3C47F21456789F0
N: [128.96.41.12]
L: p:10, a:PCMU;G728
M: sendrecv
P: PS=622, OS=31172, PR=390, OR=22561, PL=5, JI=29, LA=50

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
s=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 1296 RTP/AVP 96
a=rtpmap:96 G726-32/8000
```

If both a local and a remote connection descriptor are provided, the local connection descriptor will be the first of the two. If a connection descriptor is requested, but it does not exist for the connection audited, that connection descriptor will appear with the SDP protocol version field only.

7.3.8 RestartInProgress

The response to a RestartInProgress may include the name of another Call Agent to contact, for instance when the Call Agent redirects the endpoint to another Call Agent as in:

```
521 1204 Redirect
N: CA-1@whatever.net
```

7.4 Session description encoding

The session description is encoded in conformance with the Session Description Protocol (SDP), however, embedded clients may make certain simplifying assumptions about the session description as specified in the following. It should be noted, that session descriptions are case sensitive per IETF RFC 2327 [14].

SDP usage depends on the type of session, as specified in the "media" parameter:

- If the media is set to "audio", the session description is for an audio service.
- If the media is set to "video", the session description is for a video service.

For an audio service, the gateway will consider the information provided in SDP for the "audio" media, and for a video service the gateway will consider the information provided in SDP for the "video" media.

7.4.1 SDP audio service use

In a voice-only gateway, we only have to describe sessions that use exactly one media, audio. The parameters of SDP that are relevant for the voice based application are specified below. Embedded clients MUST support session descriptions that conform to these rules and in the following order:

- 1) The SDP profile presented below.
- 2) IETF RFC 2327 [14] (SDP: Session Description Protocol).

The SDP profile provided describes the use of the session description protocol in NCS. The general description and explanation of the individual parameters can be found in IETF RFC 2327 [14], however below we detail what values NCS endpoints need to provide for these fields (send) and what NCS endpoints should do with values supplied or not supplied for these fields (receive).

7.4.1.1 Protocol version (v=)

v= <version>
v= 0

Send: MUST be provided in accordance with IETF RFC 2327 [14] (i.e. v=0);

Receive: MUST be provided in accordance with IETF RFC 2327 [14].

7.4.1.2 Origin (o=)

The origin field consists (o=) of 6 sub-fields in IETF RFC 2327 [14]:

o= <username> <session-ID> <version> <network-type> <address-type> <address>
o= - 2987933615 2987933615 IN IP4 126.16.64.4

Username

Send: Hyphen MUST be used as username when privacy is requested.
 Hyphen SHOULD be used otherwise (see note).

NOTE: Since NCS endpoints do not know when privacy is requested, they SHOULD always use a hyphen.

Receive: This field SHOULD be ignored.

Session-ID

Send: MUST be in accordance with IETF RFC 2327 [14] for interoperability with non-IPCablecom clients.

Receive: This field SHOULD be ignored.

Version

Send: In accordance with IETF RFC 2327 [14].

Receive: This field SHOULD be ignored.

Network Type

Send: Type "IN" MUST be used.

Receive: This field SHOULD be ignored.

Address Type

Send: Type "IP4" MUST be used

Receive: This field SHOULD be ignored.

Address:

Send: MUST be in accordance with IETF RFC 2327 [14] for interoperability with non-IPCablecom clients.

Receive: This field MUST be ignored.

7.4.1.3 Session name (s=)

s= <session-name>

s= -

Send: Hyphen MUST be used as Session name.

Receive: This field MUST be ignored.

7.4.1.4 Session and media information (i=)

i= <session-description>

Send: For NCS, the field MUST NOT be used.

Receive: This field MUST be ignored.

7.4.1.5 URI (u=)

u= <URI>

Send: For NCS, the field MUST NOT be used.

Receive: This field MUST be ignored.

7.4.1.6 E-mail address and phone number (e=, p=)

e= <e-mail-address>

p= <phone-number>

Send: For NCS, the field MUST NOT be used.

Receive: This field MUST be ignored.

7.4.1.7 Connection data (c=)

The connection data consists of 3 sub-fields:

c= <network-type> <address-type> <connection-address>

c= IN IP4 10.10.111.11

Network Type:

Send: Type "IN" MUST be used.

Receive: Type "IN" MUST be present.

Address Type:

Send: Type "IP4" MUST be used

Receive: Type "IP4" MUST be present.

Connection Address:

Send: This field MUST be filled with a unicast IP address at which the application will receive the media stream. Thus a TTL value MUST NOT be present and a "number of addresses" value MUST NOT be present. The field MUST NOT be filled with a fully-qualified domain name instead of an IP address. A non-zero address specifies both the **send and receive address for the media stream(s) it covers**.

Receive: A unicast IP address or a fully qualified domain name MUST be present. A non-zero address specifies both the send and receive address for the media stream(s) it covers.

7.4.1.8 Bandwidth (b=)

b= <modifier> : <bandwidth-value>

b= AS : 64

Send: Bandwidth information is optional in SDP but it SHOULD always be included (see note 1). When an rtpmap or a non well-known codec (see note 2) is used, the bandwidth information MUST be used.

Receive: Bandwidth information SHOULD be included. If a bandwidth modifier is not included, the receiver MUST assume reasonable default bandwidth values for well-known codecs.

NOTE 1: If this field is not used, the Gate Controller might not authorize the appropriate bandwidth.

NOTE 2: A non well-known codec is a codec not defined in the codec specification J.acr.

Modifier:

Send: Type "AS" MUST be used.

Receive: Type "AS" MUST be present.

Bandwidth Value:

Send: The field MUST be filled with the Maximum Bandwidth requirement of the Media stream in kilobits per second.

Receive: The maximum bandwidth requirement of the media stream in kilobits per second MUST be present.

7.4.1.9 Time, repeat times and time zones (t=, r=, z=)

t= <start-time> <stop-time>

t= 36124033 0

r= <repeat-interval> <active-duration> <list-of-offsets-from-start-time>

z= <adjustment-time> <offset>

Send: Time MUST be present; start time MAY be zero, but SHOULD be the current time, and stop time SHOULD be zero. Repeat Times, and Time Zones SHOULD NOT be used, if they are used it should be in accordance with IETF RFC 2327 [14].

Receive: If any of these fields are present, they SHOULD be ignored.

7.4.1.10 Encryption keys

k= <method>

k= <method> : <encryption-keys>

Security services for IPCablecom are to be defined by the IPCablecom Security specification. The security services specified for RTP and RTCP do not comply with those of IETF RFC 1889 [9], IETF RFC 1890 [10], and IETF RFC 2327 [14]. In the interest of interoperability with non-IPCablecom devices, the "k" parameter will therefore not be used to convey security parameters.

Send: MUST NOT be used.

Receive: This field SHOULD be ignored.

7.4.1.11 Attributes (a=)

a= <attribute> : <value>

a= rtpmap : <payload type> <encoding name>/<clock rate> [/<encoding parameters>]

a= rtpmap : 0 PCMU / 8000

a= fmtp:*<format><format specific parameters>*

a= X-pc-codecs: <alternative 1> <alternative 2> ...

a= X-pc-secret: <method>:<encryption key>[pad]

a= X-pc-csuites-rtp: <alternative 1> <alternative 2> ...

a= X-pc-csuites-rtcp: <alternative 1> <alternative 2> ...

a= X-pc-bridge: <number-ports>

a= X-pc-nrekey: <value>

a= <attribute>

a= recvonly

a= sendrecv

a= sendonly

a=ptime

Send: One or more of the "a" attribute lines specified below MAY be included.

Receive: One or more of the "a" attribute lines specified below MAY be included and MUST be acted upon accordingly. "a" attribute lines not specified below may be present but MUST be ignored.

rtpmap:

Send: When used, the field MUST be used in accordance with IETF RFC 2327 [14]. This field MAY be used for well-known as well as non well-known codecs. The encoding names used are provided in a separate IPCablecom specification. (TS 101 909-11 [28]) On a given connection, the dynamic payload type for a given encoding method MUST be the same in the send and receive direction. Furthermore, on a given connection, once a dynamic payload type has been mapped to a given encoding method, that payload type MUST NOT subsequently be mapped to another encoding method.

Receive: When used, the field MUST be used in accordance with IETF RFC 2327 [14]. For a given connection, implementations SHOULD NOT fail if a given encoding method is mapped to different payload types in the send and receive direction, or if a given payload type is remapped.

fmtp:

Send: This field MAY be used to provide parameters specific to a particular format. For example, the field could be used to describe telephone events supported for an IETF RFC 2833 [29] format. When used, the format MUST be one of the formats specified for the media. The parameters specified are provided in a separate specification that details the usage of the format.

Receive: When used, the field MUST be used in accordance with RFC 2327 [14].

X-pc-codecs:

Send: The field contains a list of alternative codecs that the endpoint is capable of using for this connection. The list is ordered by decreasing degree of preference, i.e. the most preferred alternative codec is the first one in the list. A codec is encoded similarly to "encoding name" in rtpmap.

Receive: Conveys a list of codecs that the remote endpoint is capable of using for this connection. The codecs MUST NOT be used until signalled through a media (m=) line.

X-pc-secret:

- Send:** The field contains an end-to-end secret and (possibly) the PAD to be used for RTP and RTCP security. The secret and pad are encoded similarly to the encryption key (k=) parameter of IETF RFC 2327 [14] with the following constraints:
- The encryption key/pad **MUST NOT** contain a ciphersuite, only a passphrase.
 - The <method> specifying the encoding of the pass-phrase **MUST** be either "clear" or "base64" as defined in RFC 2045 [11], except for the maximum line length which is not specified here. The method "clear" **MUST NOT** be used if the secret or pad contains any characters that are prohibited in SDP.

The requirements for when to transmit PAD are described in IPCablecom Security, TS 101 909-11 [28]. If present, it **MUST** be separated by at least one whitespace from the secret. Pad and secret **MUST** use the same encode method.

- Receive:** Conveys the end-to-end secret and PAD to be used for RTP and RTCP security. If present, its use is as described in IPCablecom Security [28], and **MUST** be separated by at least one white space from the secret. Pad and secret **MUST** use the same encode method.

X-pc-csuites-rtp:

X-pc-csuites-rtcp:

- Send:** The field contains a list of ciphersuites that the endpoint is capable of using for this connection (respectively RTP and RTCP). The first ciphersuite listed is what the endpoint is currently expecting to use. Any remaining ciphersuites in the list represent alternatives ordered by decreasing degree of preference, i.e. the most preferred alternative ciphersuite is the second one in the list. A ciphersuite is encoded as specified below:

- ciphersuite = [AuthenticationAlgorithm] "/" [EncryptionAlgorithm]
- AuthenticationAlgorithm = 1*(ALPHA / DIGIT / "-" / "_")
- EncryptionAlgorithm = 1*(ALPHA / DIGIT / "-" / "_")

where ALPHA, and DIGIT are defined in IETF RFC 2234 [12]. Whitespaces are not allowed within a ciphersuite. The following example illustrates the use of ciphersuite:

- 62/51.

The actual list of ciphersuites to be provided in the IPCablecom Security Specification.

- Receive:** Conveys a list of ciphersuites that the remote endpoint is capable of using for this connection. Any other ciphersuite than the first in the list cannot be used until signalled through a new ciphersuite line with the desired ciphersuite listed first.

X-pc-spi-rtcp:

- Send:** The field contains the IPSEC Security Parameter Index (SPI) to be used when sending RTCP packets to the endpoint for the media stream in question. The SPI is a 32-bit identifier encoded as a string of up to 8 hex characters. The field **MUST** be supplied when RTCP security is used.

- Receive:** Conveys the IPSEC SPI to be used when sending RTCP packets over IPSEC. The field **MUST** be present when RTCP security is used.

X-pc-bridge:

- Send:** NCS endpoints **MUST NOT** use this attribute.

- Receive:** NCS endpoints **MUST** ignore this attribute if received.

X-pc-nrekey:

Send: The field contains a 16 bit integer counter for the number of rekey events. This field may be required when voice security is used. Requirements for its usage are defined in IPCablecom Security, TS 101 909-11 [28].

Receive: Conveys the number of rekey events. The field may be present when RTP security is used and its use is as defined in IPCablecom Security, TS 101 909-11 [28].

recvonly:

Send: The field MUST be used in accordance with IETF RFC 2543 [15].

Receive: The field MUST be used in accordance with IETF RFC 2543 [15].

sendrecv:

Send: The field MUST be used in accordance with IETF RFC 2543 [15].

Receive: The field MUST be used in accordance with IETF RFC 2543 [15].

sendonly:

Send: The field MUST be used in accordance with IETF RFC 2543 [15], except that the IP address and port number MUST NOT be zeroed.

Receive: The field MUST be used in accordance with IETF RFC 2543 [15].

ptime:

Send: The ptime SHOULD always be provided and when used it MUST be used in accordance with IETF RFC 2327 [14]. When an rtpmap or non well-known codec is used, the ptime MUST be provided.

Receive: The field MUST be used in accordance with IETF RFC 2327 [14]. When "ptime" is present, the MTA MUST use the ptime in the calculation of QoS reservations. If "ptime" is not present, the MTA MUST assume reasonable default values for well-known codecs.

7.4.1.12 Media Announcements (m=)

Media Announcements (m=) consists of 4 sub-fields:

M= <media> <port> <transport> <format> [,<format>]

M= audio 3456 RTP/AVP 0 97

Media:

Send: The "audio" media type MUST be used.

Receive: The type received MUST be "audio".

Port:

Send: MUST be filled in accordance with IETF RFC 2327 [14]. The port specified is the receive port, regardless of whether the stream is unidirectional or bidirectional. The sending port may be different.

Receive: MUST be used in accordance with IETF RFC 2327 [14]. The port specified is the receive port. The sending port may be different.

Transport:

Send: The transport protocol "RTP/AVP" MUST be used.

Receive: The transport protocol MUST be "RTP/AVP".

Media Formats:

Send: Appropriate media type as defined in IETF RFC 2327 [14] MUST be used.

Receive: In accordance with IETF RFC 2327 [14].

7.4.2 SDP video service use

Details on SDP use for video service are for further study.

7.5 Transmission over UDP

7.5.1 Reliable message delivery

MGCP messages are transmitted over UDP. Commands are sent to one of the IP addresses defined in the Domain Name System (DNS) for the specified endpoint or Call Agent. The responses are sent back to the source address of the command. However, it should be noted that the response may, in fact, come from another IP address than the one to which the command was sent.

When no port is provisioned for the end-point (fn 30), the commands should be sent to the default MGCP port, 2427. An end-point MUST support 4KByte datagram size on all the interfaces at the same time.

NOTE: Each endpoint may be provisioned with a separate Call Agent address and port.

MGCP messages, carried over UDP, may be subject to losses. In the absence of a timely response, commands are repeated. MGCP entities are expected to keep, in memory, a list of the responses sent to recent transactions, i.e. a list of all the responses sent over the last T_{hist} seconds, as well as a list of the transactions that are being executed currently. Transaction identifiers of incoming commands are compared to transaction identifiers of the recent responses. If a match is found, the MGCP entity does not execute the transaction, but simply repeats the response. If no match is found, the MGCP entity examines the list of currently executing transactions. If a match is found, the MGCP entity will not execute the transaction, which is simply ignored.

It is the responsibility of the requesting entity to provide suitable timeouts for all outstanding commands and to retry commands when timeouts have been exceeded. A retransmission strategy is specified in clause 7.5.2.

Furthermore, when repeated commands fail to get a response, the destination entity is assumed to be unavailable. It is the responsibility of the requesting entity to seek redundant services and/or clear existing or pending connections as specified in clause 6.4.

7.5.2 Retransmission strategy

The present document avoids specifying any static values for the retransmission timers since these values are typically network-dependent. Normally, the retransmission timers should estimate the timer by measuring the time spent between sending a command and the return of a response. Embedded clients MUST implement a retransmission strategy using exponential back-off with configurable initial and maximum retransmission timer values.

Embedded clients SHOULD use the algorithm implemented in TCP-IP, which uses two variables:

- The Average Acknowledgement Delay (AAD), estimated through an exponentially smoothed average of the observed delays.
- The Average DEVIation (ADEV), estimated through an exponentially smoothed average of the absolute value of the difference between the observed delay and the current average.

The retransmission timer, RTO, in TCP, is set to the sum of the average delay plus N times the average deviation, where N is a constant.

After any retransmission, the MGCP entity should do the following:

- It should double the estimated value of the average delay, AAD.
- It should compute a random value, uniformly distributed between 0,5 AAD and AAD.

- It should set the retransmission timer (RTO) to the minimum of:
 - the sum of that random value and N times the average deviation.
- RTO_{max} , where the default value for RTO_{max} is 4 s.

This procedure has two effects. Because it includes an exponentially increasing component, it will automatically slow down the stream of messages in case of congestion subject to the needs of real-time communication. Because it includes a random component, it will break the potential synchronization between notifications triggered by the same external event.

The initial value used for the retransmission timer is 200 ms by default and the maximum value for the retransmission timer is 4 s by default. These default values may be altered by the provisioning process.

7.6 Piggy-backing

There are cases when a Call Agent will want to send several messages at the same time to one or more endpoints in a gateway and vice versa. When several messages have to be sent in the same UDP packets, they are separated by a line of text that contains a single dot, as in for example:

```
200 2005 OK
.
DLCX 1244 aaln/2@rgw.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
```

The piggy-backed messages **MUST** be processed as if they had been received one at a time in several separate datagrams. Each message in the datagram must be processed to completion and in order starting with the first message, and each command **MUST** be responded to.

Errors encountered in a message that was piggy-backed **MUST NOT** affect any of the other messages received in that packet - each message is processed on its own.

Piggy-backing can be used to achieve two things:

- Guaranteed in-order delivery and processing of messages.
- Fate sharing of message delivery.

When piggy-backing is used to guarantee in-order delivery of messages, entities **MUST** ensure that this in-order delivery property is retained on retransmissions of the individual messages. An example of this is when multiple Notify's are sent using piggy-backing (as described in clause 6.4.3.1).

Fate sharing of message delivery ensures that either all the messages are delivered, or none of them are delivered. When piggy-backing is used to guarantee this fate sharing, entities **MUST** also ensure that this property is retained upon retransmission. For example, upon receiving a Notify from an endpoint operating in lockstep mode, the Call Agent may wish to send the response and a new NotificationRequest command in a single datagram to ensure message delivery fate sharing of the two.

7.7 Transaction identifiers and three ways handshake

Transaction identifiers are integer numbers in the range from 1 to 999,999,999. Call-agents may decide to use a specific number space for each of the gateways that they manage, or to use the same number space for all gateways that belong to some arbitrary group. Call agents may decide to share the load of managing a large gateway between several independent processes. These processes will share the same transaction number space. There are multiple possible implementations of this sharing, such as having a centralized allocation of transaction identifiers, or pre-allocating non-overlapping ranges of identifiers to different processes. The implementations **MUST** guarantee that unique transaction identifiers are allocated to all transactions that originate from any Call Agent sent to a particular gateway within a period of T_{hist} seconds. Gateways can simply detect duplicate transactions by looking at the transaction identifier only.

The Response Acknowledgement parameter can be found in any command. It carries a set of "confirmed transaction-id ranges" for final responses received - provisional responses **MUST NOT** be confirmed.

MGCP gateways may choose to delete the copies of the responses to transactions whose id is included in "confirmed transaction-id ranges" received in a message, however the fact that the transaction was executed MUST still be retained for T_{hist} seconds. Also, when a Response Acknowledgement message (see note) is received, the response that is being acknowledged by it can be deleted. Gateways should silently discard further commands from that Call Agent when the transaction-id falls within these ranges, and the response was issued less than T_{hist} seconds ago.

NOTE: As opposed to a command with a Response Acknowledgement parameter.

Let $term_{new}$ and $term_{old}$ be the endpoint-name in respectively a new command, cmd_{new} , and some old command, cmd_{old} . The transaction-ids to be confirmed in cmd_{new} SHOULD then be determined as follows:

- If $term_{new}$ does not contain any wildcards:
 - Unconfirmed responses to old commands where $term_{old}$ equals $term_{new}$.
 - Optionally, one or more unconfirmed responses where $term_{old}$ contained the "any-of" wildcard, and the endpoint-name returned in the response was $term_{new}$.
 - Optionally, one or more unconfirmed responses where $term_{old}$ contained the "all" wildcard, and $term_{new}$ is covered by the wildcard in $term_{old}$.
 - Optionally, one or more unconfirmed responses where $term_{old}$ contained the "any-of" wildcard, no endpoint-name was returned, and $term_{new}$ is covered by the wildcard in $term_{old}$.
- If $term_{new}$ contains the "all" wildcard:
 - Optionally, one or more unconfirmed responses where $term_{old}$ contained the "all" wildcard, and $term_{new}$ is covered by the wildcard in $term_{old}$.
- If $term_{new}$ contains the "any of" wildcard:
 - Optionally, one or more unconfirmed responses where $term_{old}$ contained the "all" wildcard, and $term_{new}$ is covered by the wildcard in $term_{old}$ if the "any of" wildcard in $term_{new}$ was replaced with the "all" wildcard.

A given response SHOULD NOT be confirmed in two separate messages.

The following examples illustrate the use of these rules:

- If $term_{new}$ is "aaln/1" and $term_{old}$ is "aaln/1" then the old response can be confirmed per rule 1a.
- If $term_{new}$ is "aaln/1" and $term_{old}$ is "*" then the old response can be confirmed per rule 1c.
- If $term_{new}$ is "aaln/*" and $term_{old}$ is "*" then the old response can be confirmed per rule 2a.
- If $term_{new}$ is "aaln/\$" and $term_{old}$ is "aaln/*" then the old response can be confirmed per rule 3a.

The "confirmed transaction-id ranges" values SHOULD NOT be used if more than T_{hist} seconds have elapsed since the gateway issued its last response to that Call Agent, or when a gateway resumes operation. In this situation, commands should be accepted and processed, without any test on the transaction-id.

Also, a response SHOULD NOT be confirmed if the response was received more than T_{hist} seconds ago.

Messages that confirm responses may be transmitted and received in disorder. The gateway shall retain the union of the confirmed transaction-ids received in recent commands.

7.8 Provisional responses

In some cases, transaction completion times may be significantly longer than otherwise (see note). NCS uses UDP as the transport protocol and reliability is achieved by selective time-out based retransmissions where the time-out is based on an estimate of the sum of the network roundtrip time and transaction completion time. Significant variance in the transaction completion time is therefore problematic when rapid message loss detection without excessive overhead is desired.

NOTE: For instance when resources are reserved and committed externally as part of a transaction.

In order to overcome this problem, a provisional response **MUST** therefore be issued if, and only if, the transaction completion time exceeds some small period of time. The provisional response acknowledges the receipt of the command although the outcome of the command may not yet be known, e.g. due to a pending resource reservation. As a guideline, a transaction that requires external communication to complete, e.g. network resource reservation, should issue a provisional response. Furthermore, if a duplicate CreateConnection or ModifyConnection command is received, and the transaction has not yet finished executing, a provisional response **MUST** then be sent back.

Pure transactional semantics would imply, that provisional responses should not return any other information than the fact that the transaction is currently executing, however an optimistic approach allowing some information to be returned enables a reduction in the delay that would otherwise be incurred in the system.

Provisional responses **MUST** only be sent in response to a CreateConnection or ModifyConnection command. In order to reduce the delay in the system, a connection identifier and session description **MUST** be included in the provisional response to the CreateConnection command. If a session description will be returned by the ModifyConnection command, the session description **MUST** be included in the provisional response here as well. If the transaction completes successfully, the information returned in the provisional response **MUST** be repeated in the final response. It is considered a protocol error not to repeat this information or to change any of the previously supplied information in a successful response. If the transaction fails, an error code is returned - the information returned previously is no longer valid.

A currently executing CreateConnection or ModifyConnection transaction **MUST** be cancelled if a DeleteConnection command for the endpoint is received. In that case, a response for the cancelled transaction **SHOULD** still be returned automatically, and a response for the cancelled transaction **MUST** be returned if a retransmission of the cancelled transaction is detected.

When a provisional response is received, the timeout period for the transaction in question **MUST** be set to a significantly higher value for this transaction (T_{longtran}). The purpose of this timer is primarily to detect endpoint failure. The default value of T_{longtran} is 5 s, however the provisioning process may alter this.

When the transaction finishes execution, the final response is sent and the by now obsolete provisional response is deleted. In order to ensure rapid detection of a lost final response, final responses issued after provisional responses for a transaction **MUST** be acknowledged. The endpoint **MUST** therefore include an empty "ResponseAck" parameter in those, and only those, final responses. The presence of the "ResponseAck" parameter in the final response will trigger a "Response Acknowledgement" response to be sent back to the endpoint. The "Response Acknowledgement" response will include the transaction-id of the response it acknowledges in the response header. Receipt of this "Response Acknowledgement" response is subject to the same time-out and retransmission strategies and procedures as responses to commands (see clause 6.4), i.e. the sender of the final response will retransmit it if the "Response Acknowledgement" is not received in time. The "Response Acknowledgment" response is never acknowledged.

8 Security

If unauthorized entities could use the MGCP, they would be able to set up unauthorized calls or interfere with authorized calls. Security is not provided as an integral part of MGCP. Instead MGCP assumes the existence of a lower layer providing the actual security.

Security requirements and solutions for NCS are to be provided in the IPCablecom Security Specification, which should be consulted for further information.

Annex A (normative): Event Packages

This clause defines an initial set of event packages for the various types of endpoints currently defined by IPCablecom for embedded clients. The following packages are defined for the embedded client endpoint-types listed:

Endpoint-type	Package	Package name	Default package
Analogue Access Line	Line	L	Yes
V5 LE Network Interface	ETSI	E	No
Video	For further study	For further study	For further study
ISDN BRI	For further study	For further study	For further study

Each package defines a package name for the package and event codes and definitions for each of the events in the package. In the tables of events/signals for each package, there are five columns:

Code: The package unique event code used for the event/signal.

Description: A short description of the event/signal.

Event: A check mark appears in this column if the event can be Requested by the Media Gateway Controller. Alternatively, one or more of the following symbols may appear:

"P": indicating that the event is persistent;

"S": indicating that the event is an event-state that may be audited;

"C": indicating that the event/signal may be detected/applied on a connection.

Signal: If nothing appears in this column for an event, then the event cannot be signalled on command by the Media Gateway Controller. Otherwise, the following symbols identify the type of event:

"OO": On/Off signal. The signal is turned on until commanded by the Media Gateway Controller to turn it off, and vice versa.

"TO": Timeout signal. The signal lasts for a given duration unless it is superseded by a new signal. Default time-out values are supplied. A value of zero indicates that the time-out period is infinite. The provisioning process may alter these default values.

"BR": Brief signal. The event has a short, known duration.

Additional info: Provides additional information about the event/signal, e.g. the default duration of TO signals.

Unless otherwise stated, all of the events/signals are detected/applied on endpoints and audio generated by them is not forwarded on any connection the endpoint may have. Audio generated by events/signals that are detected/applied on a connection will however be forwarded on the associated connection irrespective of the connection mode.

The "E" line package is **NOT A REPLACEMENT** for the "L" package and in fact, one absolutely must support the "L" package to create a working service for the "E" package to operate. By declaring support for the "L" with "E" line package the CA, IPAT and CM/MTAs will clearly be identifying their ability to support the European Telephony Services and Features.

Analogue Access Lines

The following packages are currently defined for Analogue Access Line endpoints. These packages apply to all endpoints:

- Line.

Package name: L

The following codes are used to identify events and signals for the "line" package for "analogue access lines":

Code	Description	Event	Signal	Additional Info
0-9, *, #, A, B, C, D	MFPB (DTMF) tones	√	BR	
bz	Busy tone	-	TO	Time-out = 30 s
cf	Confirmation tone	-	BR	
ci(ti, nu, na)	Caller Id	-	BR	"ti" denotes time, "nu" denotes number, and "na" denotes name
dl	Dial tone	-	TO	Time-out = 16 s
ft	Fax tone	√	-	
hd	Off-hook transition	P, S	-	
hf	Flash hook	P	-	
hu	On-hook transition	P, S	-	
L	MFPB (DTMF) long duration	√	-	
ld	Long duration connection	C	-	
ma	Media start	C	-	
mt	Modem tones	√	-	
mwi	Message waiting indicator	-	TO	Time-out = 16 s
oc	Operation complete	√	-	
of	Operation failure	√	-	
osi	Open interval	-	TO	default=900ms
ot	Off-hook warning tone	-	TO	Time-out = infinite
r0, r1, r2, r3, r4, r5, r6 or r7	Distinctive ringing (0..7)	-	TO	Time-out = 180 s
rg	Ringing	-	TO	Time-out = 180 s
ro	Reorder tone	-	TO	Time-out = 30 s
rs	Ringsplash	-	BR	
rt	Ring back tone	-	C, TO	Time-out = 180 s
sl	Stutter dial tone	-	TO	Time-out = 16 s
t	Timer	√	-	
TDD	Telecomm Devices for the Deaf (TDD) tones	√	-	
vmwi	Visual message waiting indicator	-	OO	
wt1, wt2, wt3, wt4	Call waiting tones	-	TO	Time-out = 12 s
X	MFPB (DTMF) tones wildcard	√	-	Matches any of the digits "0-9"

The definition of the individual events and signals are as follows:

MFPB (DTMF) tones (0-9, *, #, A, B, C, D): Detection and generation of MFPB (DTMF) signals is described in EN 300 001 [18], clause 5. It is considered an error to try and play MFPB (DTMF) tones on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone on hook).

Busy tone (bz): Station Busy is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play busy tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Confirmation tone (cf): Confirmation Tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play confirmation tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Caller Id (ci(time, number, name)): See EN 300 659-1 [17] and EN 300 659-3 [22]. Each of the three fields is optional, however each of the commas will always be included:

- The **time** parameter is coded as "MM/DD/HH/MM", where MM is a two-digit value for Month between 01 and 12, DD is a two-digit value for Day between 1 and 31, and Hour and Minute are two-digit values coded according to military local time, e.g. 00 is midnight, 01 is 1 a.m., and 13 is 1 p.m.
- The **number** parameter is coded as an ASCII character string of decimal digits that identify the calling line number. White spaces are permitted if the string is quoted, however they will be ignored.
- The **name** parameter is coded as a string of ASCII characters that identify the calling line name. White spaces, commas, and parentheses are permitted if the string is quoted.

A "P" in the number or name field is used to indicate a private number or name, and an "O" is used to indicate an unavailable number or name. The following example illustrates the use of the caller-id signal:

```
S: ci(08/14/17/26, "33 4 92 94 42 00", ETSI)
```

Dial-tone (dl): Dial Tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play dial-tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Fax tone (ft): The fax tone event is generated whenever a fax call is detected *by presence of V.21 fax preamble*. REQ2872 The fax tone event SHOULD also be generated when the T.30 CNG tone is detected. See ITU-T Recommendations T.30 and V.21 (see bibliography).

Off-hook transition (hd): See EG 201 188 [21], clause 7 Seize signal.

Flash hook (hf): See EG 201 188 [21], clause 14.2 Register recall.

On-hook transition (hu): See EG 201 188 [21], clause 8 Clear Signal. The timing for the on-hook signal is for flash response enabled.

MFPB (DTMF) Long duration (L): The "MFPB (DTMF) Long duration" is observed when a MFPB (DTMF) signal is produced for a duration longer than two seconds. In this case, the gateway will detect two successive events: first, when the signal has been recognized, the MFPB (DTMF) signal, and then, 2 s later, the long duration signal.

Long duration connection (ld): The "long duration connection" is detected when a connection has been established for more than a certain period of time. The default value is 1 hour, however this may be changed by the provisioning process.

The event may be detected on a connection. When no connection is specified, the event applies to all connections for the endpoint, regardless of when the connections are created.

Media start (ma): The media start event occurs on a connection when the first valid (see note 1) RTP media packet is received on the connection. This event can be used to synchronize a local signal, e.g. ringback, with the arrival of media from the other party.

NOTE 1: When authentication and integrity security services are used, an RTP packet is not considered valid until it has passed the security checks.

The event may be detected on a connection. When no connection is specified, the event applies to all connections for the endpoint, regardless of when the connections are created.

Modem tones (mt): Modem tone (mt): The modem tone event is generated whenever a data call is detected by presence of V.25 answer tone (ANS) with or without phase reversal or V.8 modified answer tone (ANSam) with or without phase reversal. See ITU-T Recommendation V.25 [23] and V.8 (see bibliography).

Message Waiting Indicator (mwi): Message Waiting indicator tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play message waiting indicator on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Open Switch Interval (osi): See <voiceband data transmission interface - on-hook transmission without power ringing>. See also <call processing - abandoned call scenario > and also < K-Break (as specified in ES 201 970 [30]) - an end-of-call signal consisting of a reduction in the PSTN loop current to below 1 mA for a certain period is referred to as K-break. Two times are suggested for the break:

- a) a range of 90 ms to 130 ms;
- b) a range of 250 ms to 300 ms. This is preferred for new equipment to avoid overlapping with the register recall signal>

Operation complete (oc): The operation complete event is generated when the gateway was asked to apply one or several signals of type TO on the endpoint, and one or more of those signals completed without being stopped by the detection of a requested event such as off-hook transition or dialled digit. The completion report may carry as a parameter the name of the signal that came to the end of its live time, as in:

O: L/oc(L/d1)

When the reported signal was applied on a connection, the parameter supplied will include the name of the connection as well, as in:

O: L/oc(L/rt@0A3F58)

When the operation complete event is requested, it cannot be parameterized with any event parameters. When the package name is omitted, the default package name is assumed.

The operation complete event may additionally be generated as defined in the base protocol, e.g. when an embedded ModifyConnection command completes successfully, as in note 2:

O: L/oc(B/C)

NOTE 2: Note the use of "B" here as the prefix for the parameter reported.

Operation failure (of): In general, the operation failure event may be generated when the endpoint was asked to apply one or several signals of type TO on the endpoint, and one or more of those signals failed prior to timing out. The completion report may carry as a parameter the name of the signal that failed, as in:

O: L/of(L/rg)

When the reported signal was applied on a connection, the parameter supplied will include the name of the connection as well, as in:

O: L/of(L/rt@0A3F58)

When the operation failure event is requested, event parameters can not be specified. When the package name is omitted, the default package name is assumed.

The operation failure event may additionally be generated as specified in the base protocol, e.g. when an embedded ModifyConnection command fails, as in note 2:

O: L/of(B/C(M(sendrecv(AB2354))))

Off-hook warning tone (ot): Receiver Off Hook Tone (ROH Tone) or "howler" tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play off-hook warning tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Distinctive ringing (r0, r1, r2, r3, r4, r5, r6 or r7): These power ring cadences are defined by the local administration and MAY be re-defined via provisioning.

See EG 201 188 [21] and EN 300 001 [18], clause 3. It is considered an error to try and ring a phone that is on line (off hook) and an error should consequently be returned when such attempts are made (error code 401 - phone on line (off hook)).

Ringling (rg): This power ring signal is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 3. The ringling signal may be parameterized with the signal parameter "rep" which specifies the maximum number of ringling cycles (repetitions) to apply. The following will apply the ringling signal for up to 6 ringling cycles:

```
S: rg(rep=6)
```

It is considered an error to try and ring a phone that is on line (off hook) and an error should consequently be returned when such attempts are made (error code 401 - phone on line (off hook)).

Reorder tone (ro): Re-order tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and play reorder tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Ringsplash (rs): Ringsplash, also known as "Reminder ring" is a burst of power ringling that may be applied to the physical forwarding line (when idle) to indicate that a call has been forwarded and to remind the user that a Call Forwarding subfeature is active. This signal is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 3. It is considered an error to try and ring a phone that is on line (off hook) and an error should consequently be returned when such attempts are made (error code 401 - phone on line (off hook)).

Ring back tone (rt): Audible Ring Tone is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. The ringback signal can be applied to both an endpoint and a connection.

When the ringback signal is applied to an endpoint, it is considered an error to try and play ring back tones, if the endpoint is considered off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)). When the ringback signal is applied to a connection, no such check is to be made.

Stutter Dial tone (sl): Stutter Dial Tone (also called Recall Dial Tone) is defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. The stutter dial tone signal may be parameterized with the signal parameter "del" which will specify a delay in ms to apply between the confirmation tone and the dial tone (see note 3). The following will apply stutter dial tone with a delay of 1,5 s between the confirmation tone and the dial tone:

```
S: sl(del=1500)
```

NOTE 3: This feature is needed for, e.g. Speed Dialling.

It is considered an error to try and play stutter dial tone on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

Timer (t): As described in clause 6.15, timer T is a provisionable timer that can only be cancelled by MFPB (DTMF) input. When timer T is used with the "accumulate according to digit map" action, the timer is not started until the first digit is entered, and the timer is restarted after each new digit is entered until either a digit map match or mismatch occurs. In this case, timer T functions as an inter-digit timer and takes on one of two values, T_{par} or T_{crit} . When at least one more digit is required for the digit string to match any of the patterns in the digit map, timer T takes on the value T_{par} , corresponding to partial dial timing. If a timer is all that is required to produce a match, timer T takes on the value T_{crit} corresponding to critical timing. An example use is:

```
S: dl
R: [0-9T](D)
```

When timer T is used without the "accumulate according to digit map" action, timer T takes on the value T_{crit} , and the timer is started immediately and simply cancelled (but not restarted) as soon as a digit is entered. In this case, timer T can be used as an inter-digit timer when overlap sending is used, e.g.:

```
R: [0-9](N), T(N)
```

NOTE 4: That only one of the two forms can be used at a time, since a given event can only be specified once.

The default value for T_{par} is 16 s and the default value for T_{crit} is 4 s. The provisioning process may alter both of these.

Telecomm Devices for the Deaf tones (TDD): The TDD event is generated whenever a TDD call is detected - see e.g. ITU-T recommendation V.18 (see bibliography).

Visual Message Waiting Indicator (vmwi): The transmission of the VMWI messages will conform to the requirements in EN 300 659-1 [17], clause 6.2 Data transmission not associated with ringing and EN 300 659-3 [22], clause 5.2.2 Message Waiting Indicator message. VMWI messages will only be sent from the embedded client to the attached equipment when the line is idle. If new messages arrive while the line is busy, the VMWI indicator message will be delayed until the line goes back to the idle state. The Call Agent should periodically refresh the CPE's visual indicator.

Call Waiting tone1 (wt1, ..., wt4): Call Waiting tones are defined by the local administration, and MAY be re-defined via provisioning. See EG 201 188 [21] and EN 300 001 [18], clause 1. It is considered an error to try and apply call waiting tones on a phone that is off line (on hook) and an error should consequently be returned when such attempts are made (error code 402 - phone off line (on hook)).

MFPB (DTMF) tones wildcard (X):

The MFPB (DTMF) tones wildcard matches any MFPB (DTMF) digit between 0 and 9.

Video: Event packages for video is for further study.

ISDN: Event packages for basic access ISDN is for further study.

Annex B (normative): Application of the NCS protocol to a SCN IPAT

B.1 Overview

This annex specifies an application of the NCS protocol, described in the body of the present document, to an IPAT device that is capable of emulating an Access Network to an ETSI compliant Local Exchange, (LE) that forms a part of a SCN. This annex specifies the mapping between the NCS protocol and a subset of the V5.2 protocol, see EN 300 324-1 [3] applicable for the support of SCN services to analogue telephones.

NOTE 1: This annex has been produced in response to requests of current European Cable Operators to provide telephony services over their HFC cable plants while using existing V5 Switch capacity for SCN access, as described by ECCA EuroPacketCable working group requirements document (EPC-RequDoc-V10-220501 [19]).

This annex applies to a subset of the V5 signalling protocol that relates to services provided by a 2-wire (a-b terminals), loop start, analogue POTS line.

NOTE 2: Support for additional line types is for further study.

NOTE 3: It should be recognized that while the proposed protocol enables the support of the suite of V5 SCN POTS services that due to evolving market requirements some of these services may no longer be desired or may have been discontinued within some administration boundaries. Therefore it is recommended that product compliance with the protocol in support of these services be based on manufacturers declaration, similar to practices followed with V5 PICS declarations and not on "mandated" Services compliance. In cases where a product may not support a specific service, protocol compliance should be interpreted as being able to accept the protocol interface and mitigate mismatches in service requests to product capabilities. In this way product complexity and cost can be optimized per market requirements and administration needs while maintaining protocol inter-operability.

NOTE 4: The description of the signals defined for automatic metering given by this annex and that described for a standalone metering package in annex F are the same intentionally, and should remain aligned. The equivalence of the meter pulse signals as described in this annex and that described in annex F is directly mapped; E/ps(lt=em) maps directly to am/em and E/ps(mpb) maps directly to am/mpb respectively. These signals accept the same parameter usage in both packages. Annex F provides information on the package for standalone automatic metering and if it is applied then the package shall be implemented.

NOTE 5: In the present document, ITU-T Recommendation G.711 [20] only is assumed. All other codecs must be considered as a future study.

NOTE 6: ISDN/BRI lines are for further study.

B.2 Voice gateway architecture

The reference architecture for this annex is shown in figure B.1. The IPAT provides interworking between the IPCableCom network and Local Exchange being a part of a SCN. The interface between the IPAT and the LE uses a subset of EN 300 324-1 [3] that is applicable to support of SCN services to an analogue telephone.

This mapping specified in this annex make no assumptions about the internal structure of the IPAT, however it is assumed to provide both signalling and media interworking functions.




Figure B.1: Reference model for annex B


B.3 Electrical and physical interface requirements

This proposal assumes the EN 300 324-1 [3] defined system architecture consisting of a Local Exchange (LE) and an Access Network (AN) connected via a V5 interface.

The V5 interface may have between one and sixteen 2 048 kbit/s interface as defined in EN 300 347-1 [24], EN 300 166 [25] and EN 300 167 [26].

The electrical and physical characteristics of the interface shall conform to EN 300 166 [25], 2 048 kbit/s case.

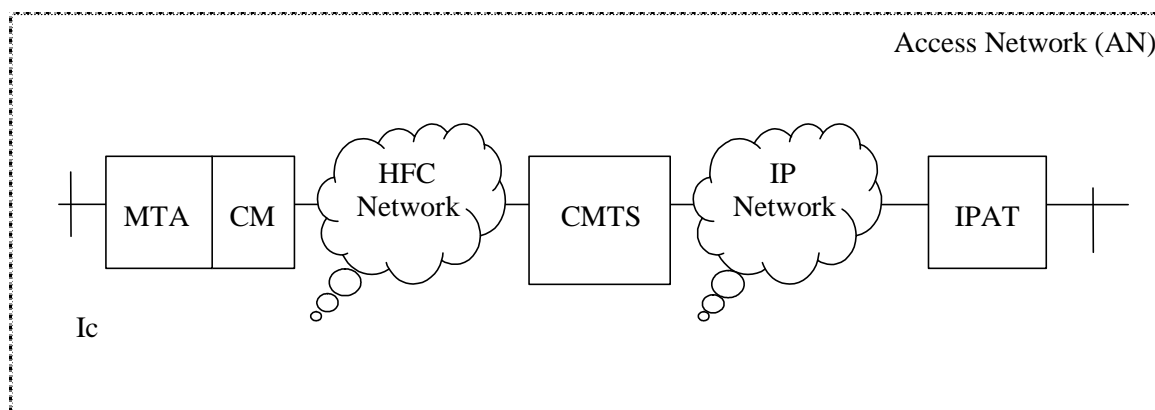
Two interface presentation alternatives are defined in EN 300 166 [25], the balanced interface pair type and the coaxial type. According to the two alternatives of interface applications shown in figure B.1, it is left to the network operator to request the interface presentation required.



NOTE: Ia = interface point at the Access Network side.
Ib = interface point at the LE side

Figure B.2

For this proposal the AN is expanded to define a PacketCable network consisting of an Internet Protocol Access Terminal (IPAT), a Cable Modem Terminal System (CMTS), a Cable Modem (CM) and a Multimedia Terminal Adapter (MTA) or an Embedded Multimedia Terminal Adapter (E-MTA).



NOTE: Ic = interface point at the user premises side.

Figure B.3

This Access Network is the synonymous to an access network utilizing a Remote Digital Terminal (RDT) in the traditional Circuit Switch architecture.

The electrical and logical definitions of the IP Network and the HFC Network are the subject of other standards activities.

This proposal assumes that these networks simply provide the transparent digital link as described in EN 300 324-1 [3]. This allows this proposal to focus on the method of providing the signalling necessary between the V5 LE to the premises interface point as defined in EN 300 324-1 [3] in support of the desired services at the user premises termination point.

For cadence ringing requests, the proposal defines an expanded range of ring cadences using a similar syntax as the NCS ringing cadence signals.

For pulsed and steady state signals, the proposal allows a PacketCable IPAT to translate a V5 protocol message received from the V5 switch to a corresponding signal request from the IPAT to the E-MTA specifying the desire signal to be applied to the premises termination point (line treatment; pulse duration, pulse period and number of repetitions, etc.). The proposal also includes a means for the IPAT to support the V5 switch requests for acknowledgements.

B.4 NCS package for V5 SCN protocol messages

This clause describes additional IPCablecom signal requests and an event request that creates the ETSI "E" Line Package which is an extension to the "L" line package described in annex A.

These signal requests and event requests map the corresponding information elements contained in a V5 SCN protocol Message Type, in a binary format, to the NCS format.

NOTE: Default values given in this annex are for the purpose of providing equipment vendors with values for initial product shipment. Provisions should be provided to allow these values to be overwritten as part of unit configuration or provisioning with alternate values per local administration requirements.

B.4.1 Cadence-ringing request

V5 "Establish" or "Signal" Message Types for "Cadence-ringing" are mapped to the NCS "SignalRequest", S: <request code>.

The signal request code for ETSI Cadence Ringing signal is **cr(x)**.

NOTE: The currently defined PacketCable Line Package NCS Ringing signal "rx" is defined with x = g, s or numbers 0 to 7 (decimal). Some of these cadences are fixed and cannot be provisioned per PacketCable guidance. V5 allows for ring cadences to range from 0 to 127, therefore the cr(x) signal request code is defined with x = 0, 127. In V5 systems the default ring cadence is cr(0) and any of the cadences can be uniquely provisioned per national norms or administration requirements.

B.4.1.1 Cadence ringing defaults and ranges

The MTA shall allow the cadence ringing values (0 through 127) to be provisioned to correspond to the LE ring cadence map per national norms or local administration requirements.

Ring Cadence Default Values are given in table B.1. All timing is in ms.

Provisioning across the range of 0 ms to 5 000 ms in steps of 50 ms is required:

Table B.1: Ring cadence default values

cr(x)	t1 - ring	t2 - idle	t3 - ring	t4 - idle	t5 - ring	t6 - idle
0	1 000	4 000	1 000	4 000	1 000	4 000
1	1 000	500	1 000	3 500	1 000	3 500
2	500	500	500	500	1 000	3 000
3	500	500	1 000	500	500	3 000
4	1 000	500	500	4 000		
5						
6						
7						
8						
....						
127						

B.4.2 Pulsed signal request

The V5 "Establish" or "Signal" Message Type " Pulsed Signal" request maps a pulsed signal request to an NCS signal request.

The signal request code for pulsed signal is **ps**.

The parameters for this signal request are:

- **lt** denotes the line treatment to be applied (Corresponds to the V5 coding of Pulse Type);
- **pd** denotes the pulse duration (length of a single pulse);
- **pr** denotes pulse repeat interval for the pulses.

The **pd** and **pr** values are optional. If no values are given the MTA shall apply pre-provisioned values in the MTA MIB per the line treatment/pulse type (lt) type code.

In addition to these parameters, the signal request may be applied with these signal request parameters:

- **rep** denotes the number of pulses (repetitions);
- **rpc** denotes the number of pulses between meter pulse reports (optional, em signal only).

Most pulsed signal requests are, in effect, TimeOut (TO) signals, in which the timeout value may be determined as:

- $to = pr * rep$.

The IPAT need not include the timeout parameter in the signal request if the default timeout value is adequate for the Signal Request being requested. This default must be provisioned in both the MTA and the IPAT.

The IPAT SHOULD include the timeout value if the product of $pr*rep$ is significantly less than 180 s, and the IPAT MUST include the timeout value if the product of $pr*rep$ is greater than 180 s.

The "enable metering pulse generation" (**em**) and "burst metering pulse generation" (**mpb**) signals are defined as on/off (OO) and brief (BR) signals, respectively. The number of pulses (**rep**) is not applicable to the **em** signal request. Rather, the **em** signal only may include the report pulse count (**rpc**) parameter. The number of pulses parameter is required for the **mpb** signal request.

B.4.2.1 Line treatment encoding

Table B.2 describes encoding for the line treatments that may be applied, along with signal type and parameter applicability. Parameters may be Optional (O), Mandatory (M) or Forbidden (F).

Table B.2: Line treatment encoding

lt Code	Description	Signal Type	pd	pr	rep (note)	rpc
ir	Initial Ring	TO	O	O	O	F
lc	Pulsed loop closed	TO	O	O	O	F
lo	Pulsed loop open	TO	O	O	O	F
em	(Enable) metering pulse generation	OO	F	O	F	O
mpb	Metering pulse burst generation	BR	O	O	O	F
nb	Pulsed no battery	TO	O	O	O	F
np	Pulsed normal polarity	TO	O	O	O	F
rb	Pulsed reduced battery	TO	O	O	O	F
rp	Pulsed reversed polarity	TO	O	O	O	F
NOTE:	"rep" parameter is MANDATORY if value is provided by the V5 LE interface. The assignment of OPTIONAL in this field is in recognition of the use of default values (table B.3) in support of Call Agent or Softswitch architectures.					

B.4.2.2 Line treatment defaults and ranges

Table B.3 describes defaults and parameter ranges for the line treatments in table B.1. Timing Values are in ms.

Table B.3: Line treatment defaults and ranges

It Code	Description	Frequency (tolerance)	Amplitude (min-max, steps)	pd (min-max, steps)	pr (min-max, steps)	rep (min-max, steps)
ir	Initial Ring	25 Hz (±1 Hz)	Full	200 (0 - 5 000, 50)	200 (0 - 5 000, 50)	1 (1 - 5, 1)
lc	Pulsed loop closed	null	null	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
lo	Pulsed loop open	null	null	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
em	(Enable) metering pulse generation	16 KHz	-13,5 dBm1 (-25 to +15,2 dB)	150 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	null
mpb	Metering pulse burst generation	16 KHz	-13,5 dBm1 (-25 to +15,2 dB)	150 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
nb	Pulsed no battery	Null	0	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
np	Pulsed normal polarity	Null	1	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
rb	Pulsed reduced battery	Null	1	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)
rp	Pulsed reversed polarity	Null	0	200 (0 - 5 000, 10)	1 000 (0 - 5 000, 10)	1 (1 - 50, 1)

NOTE: Meter Pulse Amplitude is specified in dBm across a-b terminals terminated in the reference termination impedance per national norms.

B.4.2.3 Requested events

The following events may be requested for pulsed signals, by inclusion in the requested events (R: parameter list in the notification request:

- **oc** denotes that operation completion should be notified;
- **of** denotes that operation failure should be notified;
- **pc** denotes that pulse completion should be notified.

B.4.2.4 Pulse encoding

The IPAT must map V5 enumerated pulse type and duration coding to NCS line treatment types and durations in ms per provisioning tables as defined by the LE or the local administration.

B.4.2.4.1 Pulse duration encoding

The pulse duration is specified in ms, using the **pd** parameter. For example, a 200 ms pulse is specified by:

pd=200

Pulse duration is *optional*. If not provided by the requesting entity, the MTA SHOULD apply a provisioned or internally defaulted value, based on the line treatment (It) parameter (table B.3).

B.4.2.4.2 Pulse period encoding

The pulse period is specified in ms, using the **pr** parameter. For example, a 1 s period is specified by:

pr=1 000

Thus, for example, a 50 % duty cycle, 1 s periodic pulse is specified by:

pd=500, pr=1 000

Pulse period is *optional*. If not provided by the requesting entity, the MTA SHOULD apply a provisioned or internally defaulted value, based on the line treatment (lt) parameter (table B.3).

B.4.2.5 Pulse completion event coding

The pulse completion event is reported by the MTA when requested in the first Signal Request by the IPAT and when each requested pulse is completed. This event is notified for each completed pulse for the duration of the signal request, without requiring additional notification requests from the IPAT. Detection of this event does not affect continued application of pulses by the MTA.

The event request code for pulse complete is **pc**, and is included in the signal request, similar to the operation complete **oc** event.

B.4.2.6 Metering pulse report coding

The metering pulse report event is reported by the MTA when requested in an enable metering pulse generation signal request with non-zero report pulse count (**rpc**) parameter. This event is notified each time the MTA's metering pulse count reaches the report pulse count. Generation of the event resets the MTA's metering pulse count to zero. The count does not include pulses generated by any metering pulse burst (**mpb**) signal requests. Generation of the event does not affect continued generation of metering pulses and subsequent metering pulse report event notification. The IPAT does not need to send a new notification request.

The event code for the metering pulse report is **mpr**. The notification includes the count.

EXAMPLE: O: mpr(10).

B.4.2.7 V5 suppression indicator

The V5 suppression indicator is used both in the Pulsed-Signal IE and in the Enable-Metering IE. It allows the LE to indicate to the Access Network whether the ongoing pulsed signal shall be suppressed.

The suppression indicator shall be used to indicate whether the pulse generation shall be stopped in a network if the line conditions change, if a new SIGNAL message is received from the LE or if either occurs. This is especially important for meter pulses in some networks, where metering pulses are not sent after the call is cleared, this could be used to suppress metering pulses after the call has been cleared.

In other networks it is essential that the meter pulses are sent out regardless of either a change in line state due to messages from the LE or changes due to the TE.

The coding of suppression indicator is:

- 00 No suppression;
- 01 Suppression allowed by pre-defined V5.1 SIGNAL message from LE;
- 10 Suppression allowed by pre-defined line signal from TE;
- 11 Suppression allowed by pre-defined V5.1 SIGNAL message from LE or pre-defined line signal from TE.

The signal suppression option does not map efficiently to the NCS protocol. For example, to apply a signal request with "no suppression," the signal must be defined as a "brief" signal; to apply a signal with "suppression allowed by pre-defined signal from TE" requires that the signal must be defined as a "timeout" signal. For the purposes of V5-to-NCS inter-working, the NCS behaviour is accepted, and the signals are defined based on assumptions of normal usage.

To resolve this conflict with NCS the IPAT must "bridge" the V5 protocol to NCS by accepting the V5 Suppression Indication and then executing the appropriate set of NCS messages to achieve the desired effect.

B.4.2.7.1 No suppression

Upon receipt of the V5 "00" code, the IPAT shall generate the associated line treatment NCS message to the MTA. The MTA shall execute the associated line treatment as defined in this annex regardless of changes in the line state or additional signal messages from the LE-IPAT.

B.4.2.7.2 Suppression by pre-defined V5 signal message

For this case the IPAT must be pre-provisioned with the associated V5 SIGNAL message (e.g. far end "on hook").

Upon receipt of the V5 "01" code, the IPAT shall begin monitoring for the pre-provisioned V5 SIGNAL message.

The MTA shall execute the associated line treatment as defined in this annex.

Upon receipt of the pre-provisioned V5 SIGNAL message, the IPAT shall issue the associated pulsed signal cancellation message (see clause B.4.5) to the MTA.

The MTA shall respond to the associated pulsed signal cancellation message as defined in this annex.

B.4.2.7.3 Suppression by pre-defined line signal from TE

For this case the IPAT must be pre-provisioned with the associated NCS line treatment signal message (e.g. "on hook").

Upon receipt of the V5 "10" code, the IPAT shall begin monitoring for the pre-provisioned NCS line treatment signal message from the MTA. The MTA shall execute the associated line treatment message as defined by NCS protocols (e.g. "on hook").

Upon receipt of the pre-provisioned NCS line treatment message, the IPAT shall issue the associated pulsed signal cancellation message (see clause B.4.5) to the MTA.

The MTA shall respond to the associated pulsed signal cancellation message as defined in this annex.

B.4.2.7.4 Suppression by pre-defined V5 SIGNAL message from LE or pre-defined line signal from TE

For this case the IPAT must be pre-provisioned with an associated V5 SIGNAL message **AND** the associated NCS line treatment signal message (e.g. far end "on hook" AND TE "on hook").

Upon receipt of the V5 "11" code, the IPAT shall begin monitoring for the pre-provisioned V5 SIGNAL message and shall begin monitoring for the pre-provisioned NCS line treatment signal message from the MTA.

If presented to the MTA the MTA shall execute the associated line treatment message as defined by NCS protocols (e.g. "off hook").

Upon receipt of the pre-provisioned V5 SIGNAL message **OR** the NCS line treatment message from the MTA, the IPAT shall issue the associated pulsed signal cancellation message (see clause B.4.5) to the MTA.

The MTA shall respond to the associated pulsed signal cancellation message as defined in this annex.

B.4.2.8 Repetition indicator

The repetition indicator is only used in the V5 Enable-Metering IE. It is sent in the direction of LE to Access Network with a reporting pulse count to instruct the Access Network whether to continue or cease application of automatic metering pulses when the number specified in **reporting pulse count** have been applied.

Coding of repetition indicator:

- **00 Cease to apply pulses after number specified by reporting pulse count have been applied;**
- 11 Continue to apply pulses at same rate until the call is disconnected or receipt of new instructions from LE;
- 01 Reserved for ETSI use;

- 10 Reserved for ETSI use.

The default behaviour for the **em** line treatment provides for the signal to be applied as an on/off signal until discontinued by the IPAT. The IPAT can obtain the behaviour of discontinuing the pulses once the reporting pulse count has been reached by including an embedded notification request to turn off the **em** signal (see clause B.4.5).

B.4.3 Pulse repetition encoding

The IPAT maps the V5 interface pulse repetition count directly to the existing NCS repetition (*rep*) parameter.

This parameter must be provided in accordance with table B.2. There is no default value for pulse repetitions.

NOTE: Per V5 guidance, a rep value of "0" is invalid. If the IPAT receives a request from the V5 LE with a rep value = 0, or missing, the IPAT shall substitute a rep value of "1".

In the V5 Pulsed-Signal IE, the "number of pulses" field is a 5-Bit field. The range of permitted values is 1 through 31. In the V5 Enable-Metering IE, the combination of "repetition indicator = 00" and "reporting pulse count" fields also allow for the specification of a limited "number of pulses". Reporting pulse count is a 12-Bit field, accounting for a valid range from 1 through 4 095. While the pulse repetition value may only be in the range 1..31 from a V5 interface, pulse repetitions may be specified over the full range 1 to 4 095.

B.4.4 Parameter usage

All of the parameters described for the pulsed signal request apply to all of the described line treatments.

The IPAT must supply values for the pulse duration, the pulse repetition interval, and the number of repetitions.

To account for national variation for meter pulse, the frequency and amplitude are provisioned to the MTA because neither is provided in the message from the V5 interface. The IPAT must determine the pulse repetition interval from the V5 interface message rate type and provide the interval time (ms) to the MTA in the signal request.

In V5-2000, the Enable-Metering information element has a rate type field. This is an enumeration type. The IPAT must translate the different enum values to corresponding ms values, based on its provisioning, depending on the local administration.

The IPAT may use the pulse repetition interval and signal repetition parameter to generate a fixed number of pulses to the subscriber's line.

B.4.5 Pulsed signal cancellation

Most pulsed signals, being timeout signals, are terminated when any requested event is detected, *except for pulse completion (pc)*.

In addition, the LE may terminate all active pulsed signals at any time by sending an empty Signal Request.

Since the LE may apply multiple pulsed signals to a subscriber's line simultaneously, (for example, meter pulse is being generated and another line treatment is applied), the IPAT may terminate an on/off line treatment with a treatment specific command syntax. An example to terminate the applied meter pulse would be:

- S: E/ps(em(-))

B.4.6 Pulsed completion event

The pulsed completion event is reported by the MTA to the IPAT when each requested pulse is completed.

The event request code for pulse complete is **pc**.

B.4.7 Pulsed signal failure event

The pulsed signal failure event is reported by the MTA to the IPAT when any pulsed signal request fails to complete, if operation failure "of" has been included in the list of requested events. A pulsed signal request may fail for any reason that any other signal request might fail.

B.4.8 Steady-signal request

The V5 "Establish" Steady-Signal request maps a Steadysignal request to an NCS signal request.

The signal request code for steady signal is **ss**.

The parameters for this signal request are:

- **lt** denotes the line treatment to be applied (Corresponds to the V5 coding of Steady-signal Type).

This treatment is maintained until the V5 LE directs for a new treatment.

B.4.8.1 Line treatment encoding

Line treatments are encoded using the following code words:

Table B.4: Steady signal request encoding

lt Code	Description
fb	normal (full) battery
lc	loop closed
lo	loop open
nb	no battery
np	normal polarity
rb	reduced battery
rp	reversed polarity

B.4.8.2 Line treatment provisioning

There is no provisioning required in that these are line states without quantitative values (timing, frequency or amplitude).

B.4.9 Metering pulse generation

On receiving an "enable metering pulse generation" **ps(lt=em(+))** signal request the MTA shall apply the first metering pulse to the termination immediately, and then apply subsequent metering pulses at intervals as specified by the value of the pulse repetition interval parameter **pr**, if supplied in the signal request, or the provisioned value.

The MTA shall continue to generate metering pulses until it receives a "disable metering pulse generation" **ps(lt=em(-))** signal request, or an empty signal request list.

A metering pulse burst signal **ps(lt=mpb)** request may be included in a signal request that also enables metering pulse generation, for example, to apply an initial charge to a call. When this occurs, the MTA shall apply the metering pulse burst to the endpoint completely, and then start generating normal metering pulses.

Because the metering pulse burst signal is a brief signal type, all pulses specified to for the request (**rep=n**) are applied, even if the subscriber goes on-hook during the burst.

A metering pulse burst signal request may occur during a call in progress, for example to take account of a chargeable subscriber action. When this occurs, the MTA shall suspend normal metering pulse generation, and apply the metering pulse burst signal request. The MTA then shall resume normal metering pulse generation without requiring a new "enable metering pulse generation" request from the IPAT. The IPAT must account for any normal metering pulses missed during the burst by including the missed pulses in the burst count.

The IPAT may optionally include a report pulse count (**rpc**) parameter with the enabling metering pulse generation (**em**) signal request. When this parameter is non-zero (rpc=n, where n=1 to x), the MTA generates meter pulse reports, in the form of notifications, each time its pulse count reaches the rpc value. Generation of the event notification resets an rpc counter so that a report will be generated each time the rpc "n" value is reached. This count does not include any metering pulses generated by metering pulse burst (**mpb**) signal requests.

B.5 Provisioning configurations

B.5.1 MTA

The MTA shall be provisioned with electrical parameters for each of the Line Treatments. When appropriate, these parameters include, amplitude, frequency, minimum pulse widths and maximum rep rate (minimum inter-pulse timing). See table B.1 through table B.3 for details. These parameters are to be used unless line treatment specific values are provided by the V5 interface messages.

B.5.2 IPAT

Shall be provisioned with V5 pulse type and duration coding mapping to NCS pulse type and pulse duration timing in ms. This provisioning must be consistent with the LE provisioning and the local administration guidance.

B.6 ETSI line package support

B.6.1 NCS audit

The NCS Audit Endpoint (AUPEP) command allows the MTA to report signals that it supports.

In response to an AUPEP, an MTA that supports any of the signalling requests listed in this annex must report support of this "ETSI" package (designated with an "E" code).

An example of an audit exchange:

```
AUPEP 1232 aaln/1@rgw.mso.net
F: A
```

MTA responds:

```
200 1232 OK
A: a:PCMU,
p:30-90,
v:L;E,
m:sendonly;recvonly;sendrecv;inactive,
DQ-GI,SC-ST, SC-RTP: 00/51;03
```

The important line for packages is the "v:L;E" which indicates support for the NCS Line Package (L) and for the ETSI line package (E).

B.6.2 Unsupported signals - PICS declaration

This is an indication of a device platform limitation (hardware or software) and is not an error condition.

Product vendors must reflect any unsupported signals listed in this annex in the product PICS declaration.

NCS provides a messaging facility where as, if the device cannot support the requested signal type, the device shall return an "unsupported signal" response (513 code).

EXAMPLE 1: CMS->MTA (requesting a metering pulse burst):

RQNT 9915 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0

X: 2255

S: E/ps(lt=mpb, pd=500, pr=1 000, rep=5)

R: oc, hu, hf

MTA->CMS (rejecting the request):

513 9915 Unsupported Signal in Signal Request.

EXAMPLE 2: CMS->MTA (requesting metering enable, using provisioned defaults):

RQNT 9915 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0

X: 2255

S: E/ps(lt=em(+))

R: E/pc, hu, hf

MTA->CMS (rejecting the request):

513 9915 Unsupported Signal in Signal Request.

B.7 Call flow examples

B.7.1 Cadence ringing

B.7.1.1 Cadence ringing call flow for basic ring cadence

This flow illustrates a request for the application of a simple ring cadence

- 1) The V5 LE includes a cadence-ringing pulsed signal request in a message to the IPAT.
- 2) The IPAT converts the binary coded cadence-ring to a decimal value between 0 and 127.
- 3) Assuming the cadence-ring value is converted to a decimal number "0":
 - RQNT 500 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
 - S: E/cr(0).
- 4) The MTA acknowledges the signal request:
 - 200 500 OK.
- 5) The MTA looks up in its provisioned ring table for the cr(0) definition of ring frequency and ring cadence and applies it to the a-b terminals for the aaln/1 line presence on the MTA.

This cadence continues until the MTA detects Off Hook at which time it begins the normal NCS connect sequence or until the IPAT signals a disconnect message.

B.7.1.2 Cadence ringing - Ring splash followed by a ring cadence

This Call flow demonstrates the use of a pulsed signal "initial ring" type followed by a ring cadence to provide a "ring splash" followed by a ring cadence.

- 1) The V5 LE supplies an "initial ring" pulsed signal type request with a pulse duration type in a message to the IPAT.

- 2) The IPAT converts the "initial ring" type to the NCS It type ir, with the pulse duration value and requests operation complete notification:
 - RQNT 510 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 000691
S: E/ps (lt=ir, pd = 200, rep=1)
R: oc.
- 6) The MTA acknowledges the signal request.
- 7) 200 691 OK.
- 8) The MTA looks up in its provisioned ring table for the ir definition of initial ring frequency and initial ring duration (pd = 200 results in a ring burst of 200 ms) and applies it to the a-b terminals for the aaln/1 line presence on the MTA.
- 9) Upon completion of the initial ring the MTA responds with an operation complete message:
 - a) NTFY 1298 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 691
O: oc(E/ps(ir)).

NOTE: Since the "E" line package is optional, it is necessary to use the "E/" package name notation to explicitly identify the signal as being a signal from the "E" package definition.

- 7) The IPAT signals the V5 LE that the pulse is complete.
- 8) The V5 LE includes a cadence-ringing pulsed signal request in a message to the IPAT.
- 9) The IPAT converts the binary coded cadence-ringing to a decimal value between 0 and 127.
- 10) Assuming the cadence-ringing value is converted to a decimal number "0".
- 11) RQNT 520 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 699
S: E/cr(0).
- 12) The MTA acknowledges the signal request:
 - 200 520 OK.
- 13) The MTA looks up in its provisioned ring table for the cr(0) definition of ring frequency and ring cadence and applies it to the a-b terminals for the aaln/1 line presence on the MTA.

This cadence continues until the MTA detects Off Hook at which time it begins the normal NCS connect sequence or until the IPAT signals a disconnect message.

B.7.1.3 Cadence ringing - Ring splash followed by "on hook" data, then ring cadence

This flow illustrates an "on hook" data transmission associated with ringing (CLID).

A ring burst preceding V5 LE generated FSK signalling tones followed by the application of a ring cadence.

- 1) The V5 LE supplies an "initial ring" pulsed signal type request with a pulse duration type in a message to the IPAT.
- 2) The IPAT converts the "initial ring" type to the NCS It type ir, with the pulse duration value and requests operation complete notification:
 - RQNT 530 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 777
S: E/ps (lt=ir, pd = 200, rep=1)
R: oc.

- 3) The MTA acknowledges the signal request:
 - 200 530 OK.
- 4) The MTA looks up in its provisioned ring table for the ir definition of initial ring frequency and initial ring duration (pd = 200 results in a ring burst of 200 ms) and applies it to the a-b terminals for the aaln/1 line presence on the MTA.
- 5) Upon completion of the initial ring the MTA responds with an operation complete message.
- 6) NTFY 1298 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 777
O: oc(E/ps(ir)).

NOTE: Since the "E" line package is optional, it is necessary to use the "E/" package name notation to explicitly identify the signal as being a signal from the "E" package definition.

- 7) The IPAT signals the V5 LE that the pulse is complete.
- 8) The V5 LE then generates the FSK tones in band to the aaln/1 termination.
- 9) The MTA plays through the in band FSK tones to the aaln/1 analogue POTS line.
- 10) The V5 LE times from the end of the FSK tone a 200 ms delay (to meet the minimum requirements of EN 300 659-1 [17]) and then generates a cadence-ringing pulsed signal request in a message to the IPAT.
- 11) The IPAT converts the binary coded cadence-ring to a decimal value between 0 and 127.
- 12) Assuming the cadence-ring value is converted to a decimal number "0".
- 13) RQNT 540 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 778
S: E/cr(0).
- 14) The MTA acknowledges the signal request:
 - 200 540 OK.
- 15) The MTA looks up in its provisioned ring table for the cr(0) definition of ring frequency and ring cadence and applies it to the a-b terminals for the aaln/1 line presence on the MTA.

This cadence continues until the MTA detects Off Hook at which time it begins the normal NCS connect sequence or until the IPAT signals a disconnect message.

B.7.2 Pulsed signal request

B.7.2.1 Pulse signal request for one loop open pulse

- 1) The V5 LE includes a loop open pulsed signal request in a message to the IPAT.
- 2) The IPAT converts the binary coded V5 message and determines the line treatment and pulse duration from parameters supplied by the switch, and generates an appropriate NCS signal request.

```
RQNT 525 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 795
S: E/ps(lt=10, pd=200, rep=1)
```

- 3) The MTA acknowledges the signal request.
 - 200 525 OK
- 4) The MTA applies a 200 ms open loop to the subscriber's access line.

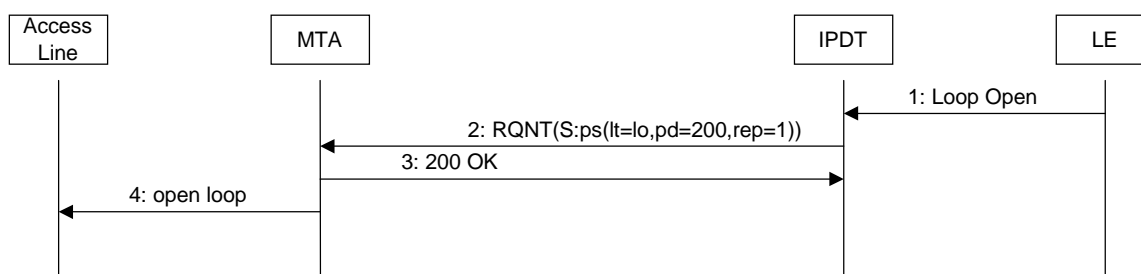


Figure B.4: Pulsed signal request

B.7.2.2 Pulsed signal with start acknowledgement

This call flow illustrates a pulsed signal request with multiple pulses, and in which the switch has requested acknowledgement when the signal application to the subscriber's access line starts.

- 1) The V5 LE request an open loop with multiple pulses and start acknowledgement.
- 2) The IPAT converts the binary coded V5 message and determines the line treatment, pulse duration and pulse period from parameters supplied by the switch, and generates an appropriate NCS signal request, including the number of pulse repetitions supplied by the V5 LE. The IPAT must "remember" that the switch has requested signal start acknowledgement

```
RQNT 525 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 919
S: E/ps(lt=lo, pd=200, pr=1 000, rep=3).
```

- 3) The MTA acknowledges the signal request.
- 4) 200 525 OK.
- 5) The IPAT sends acknowledgement to the V5 LE.
- 6) The MTA starts applying the open loop pulses to the subscriber's access line.




Figure B.5: Pulsed signal with start acknowledgement

B.7.2.3 Pulsed signal with completion acknowledgement

This call flow illustrates a pulsed signal request in which the V5 LE has requested acknowledgement after all pulses have been applied.

- 1) The V5 LE requests an open loop with multiple pulses and completion acknowledgement.
- 2) The IPAT converts the binary coded V5 message and determines the line treatment and pulse duration from parameters supplied by the switch, and generates an appropriate NCS signal request, including the number of pulse repetitions supplied by the V5 LE. Because the V5 LE also requested completion acknowledgement, the IPAT includes the operation complete parameter in the signal request. For the sake of this example, assume also that the V5 LE requested start acknowledgement

```
RQNT 525 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 942
S: E/ps(lt=lo, pd=200, pr=1 000, rep=3)
R: oc.
```

- 3) The MTA acknowledges the signal request:
 - 200 525 OK.
- 4) The MTA starts applying the requested pulses to the line.
- 5) 2nd pulse.
- 6) 3rd pulse.
- 7) With the last pulse completed, the MTA notifies the IPAT that the operation is complete:

```
NTFY 1298 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 942
O: oc(E/ps(lo)).
```

Note that this assumes an "ETSI Line Package" designated by the name "E". The package name could be omitted if this package is the default package.

- 8) The IPAT sends the requested acknowledgement to the V5 LE.
- 9) The IPAT acknowledges the event notification to the MTA.

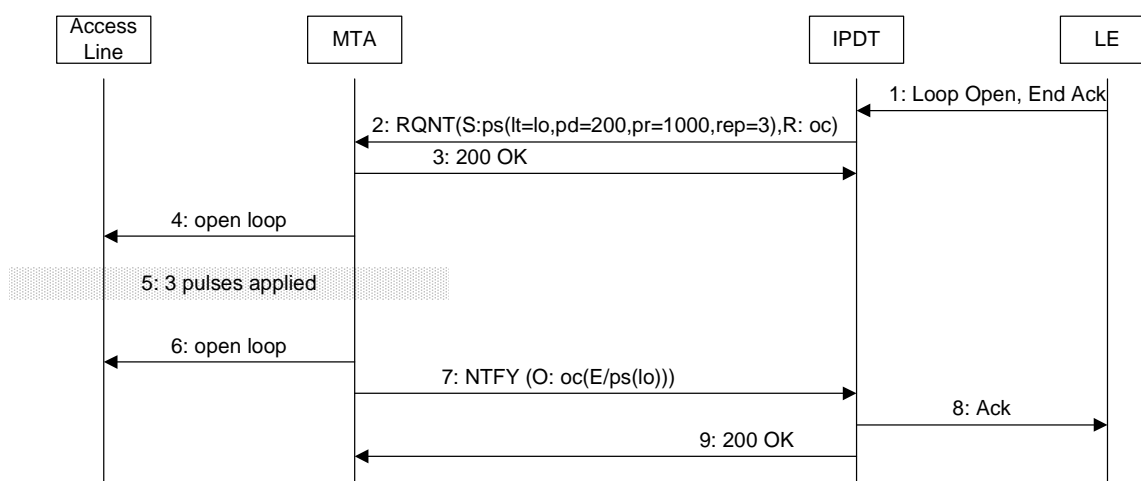


Figure B.6: Pulsed signal with completion acknowledgement

B.7.2.4 Pulsed signal with pulse acknowledgement

This call flow illustrates a pulsed signal request in which the V5 LE has requested acknowledgement after each pulse is applied.

- 1) The V5 LE requests an open loop with multiple pulses and pulse acknowledgement.
- 2) The IPAT converts the binary coded V5 message and determines the line treatment and pulse duration from parameters supplied by the switch, and generates an appropriate NCS signal request, including the number of pulse repetitions supplied by the V5 LE. Because the V5 LE also requested pulse acknowledgement, the IPAT includes an embedded signal request for the pc signal:

```
RQNT 525 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 1111
S: E/ps(lt=10, pd=200, pr=1 000, rep=3)
R: E/pc.
```

- 3) The MTA acknowledges the signal request:
 - 200 525 OK.
- 4) The MTA applies the first pulse to the subscriber's access line.
- 5) When the pulse completes, the MTA sends an event notification to the IPAT:

```
NTFY 3981 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 1111
O: E/pc(1t).
```

- 6) The IPAT sends the pulse acknowledgement to the V5 LE.
- 7) The IPAT acknowledges the event notification. The IPAT does not need to send a new notification request for pulse completion. This request remains in effect until metering pulse generation is completed.
- 8) The MTA continues to transmit pulses and notify pulse completions.




Figure B.7: Pulsed Signal with Pulse Acknowledgement

B.7.2.5 Pulsed signal - Meter pulse with pulse acknowledgement

This call flow illustrates a pulsed signal request in which the V5 LE has requested application of meter pulse with acknowledgement after each pulse is applied. The Meter Pulse frequency has been provisioned to the MTA.

- 1) The V5 LE requests enable metering pulse generation and pulse acknowledgement.
- 2) The IPAT converts the binary coded V5 message and generates an appropriate NCS signal request. Because the V5 LE also requested pulse acknowledgement, the IPAT includes the pc parameter with the signal request:

```
RQNT 535 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 2345
S: E/ps(lt=em(+))
R: E/pc.
```

- 3) The MTA acknowledges the signal request:
 - 200 535 OK.
- 4) The MTA refers to its provisioning table to determine the meter pulse frequency, amplitude and default timings, and applies the first meter pulse to the subscriber's access line.
- 5) When the pulse completes, the MTA sends an event notification to the IPAT:

```
NTFY 3981 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 535
O: pc(em).
```

- 6) The IPAT sends the pulse acknowledgement to the V5 LE.
- 7) The IPAT acknowledges the event notification.
- 8) The MTA continues to transmit pulses and notify pulse completions until the V5 LE discontinues metering pulse generations:

```
RQNT 599 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
S: E/ps(lt=em(-)).
```

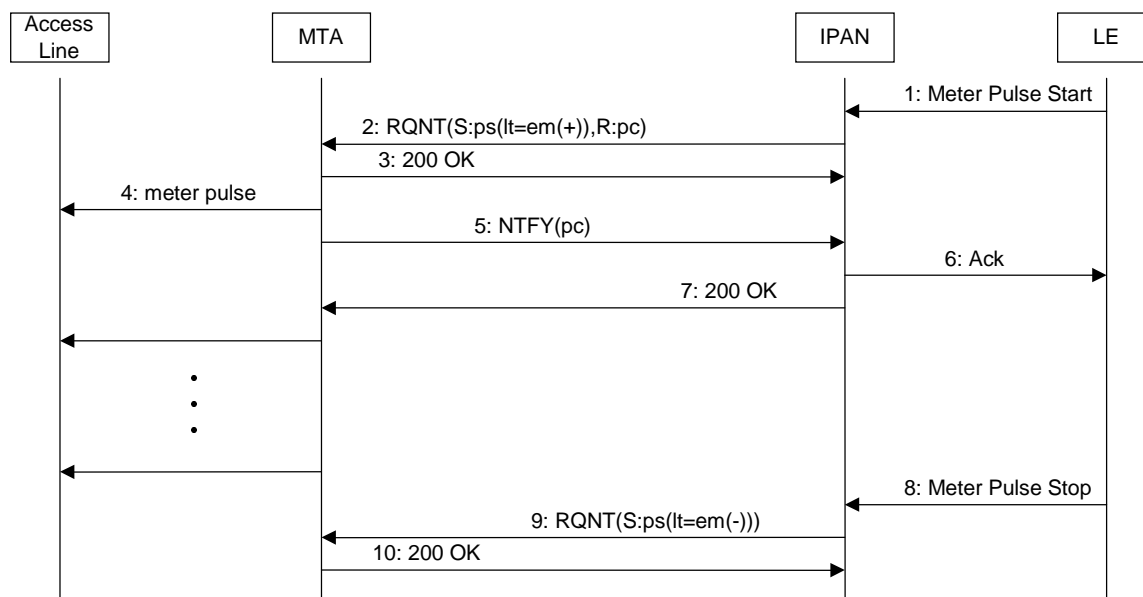


Figure B.8: Metering with pulse acknowledgement

B.7.2.6 Pulsed signal - Meter pulse with pulse acknowledgement with tariff change

This call flow illustrates a pulsed signal request in which the V5 LE has requested application of meter pulse with acknowledgement. After several pulses in the first string are applied, a tariff change is invoked. The Meter Pulse frequency has been provisioned to the MTA.

- 1) The V5 LE request application of meter pulse with multiple pulses and pulse acknowledgement.
- 2) The IPAT converts the binary coded V5 message and determines the line treatment and pulse duration from parameters supplied by the switch and generates an appropriate NCS signal request, including the number of pulse repetitions supplied by the V5 LE. Because the V5 LE also requested pulse acknowledgement, the IPAT includes an embedded signal request for the pc signal:

```
RQNT 545 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 3579
S: E/ps(lt=em(+), pd=150, pr=1 000)
R: E/pc.
```

- 3) The MTA acknowledges the signal request:
 - 200 545 OK.
- 4) The MTA refers to its provisioning table to determine the meter pulse frequency, amplitude and default timings (minimum allowed values).
- 5) The IPAT relays the start acknowledgement to the V5 LE. Cannot have both start ack and per-pulse ack.
- 6) The MTA applies the first meter pulse to the subscriber's access line.
- 7) When the pulse completes, the MTA sends an event notification to the IPAT:


```
NTFY 3981 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 3579
O: pc(em).
```
- 8) The IPAT sends the pulse acknowledgement to the V5 LE.
- 9) The IPAT acknowledges the event notification. The IPAT does not need to send a new notification request for pulse completion. This request remains in effect until metering pulse generation is completed.
- 10) The MTA continues to transmit pulses and notify pulse completions.

As the result in a change in the call state (e.g. start up of three way call) the LE determines that a new tariff is to be applied. Based on the new tariff, the LE determines a new meter pulse rate.

- 11) The V5 LE request application of meter pulse with a new multiple pulse count and start acknowledgement.
- 12) The IPAT converts the new binary coded V5 message and determines the line treatment and pulse duration from parameters supplied by the switch and generates an appropriate NCS signal request, including the number of pulse repetitions supplied by the V5 LE. Because the V5 LE also requested pulse acknowledgement, the IPAT includes an embedded signal request for the pc signal. For the sake of this example, assume also that the V5 LE requested start acknowledgement:

```
RQNT 547 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 3581
S: E/ps(lt=em(+), pd=150, pr=500)
R: E/pc.
```

- 13) The MTA acknowledges the signal request:
 - 200 547 OK.

- 14) The MTA refers to its provisioning table to determine the meter pulse frequency, amplitude and default timings (minimum allowed values).
- 15) The IPAT relays the start acknowledgement to the V5 LE.
- 16) The MTA applies the first new meter pulse to the subscriber's access line with the new pulse rate.
- 17) When the pulse completes, the MTA sends an event notification to the IPAT:


```

NTFY 791 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 3581
O: pc(em).

```
- 18) The IPAT sends the pulse acknowledgement to the V5 LE.
- 19) The IPAT acknowledges the event notification.
- 20) The MTA continues to transmit pulses and notify pulse completions.

B.7.3 Fixed meter pulse application, completed

This call flow illustrates meter pulse application with operation complete notification.

- 1) The LE requests application of twenty-five (25) meter pulses to the subscriber's access line, with pulse duration of 150 ms and repetition interval of 2 000 ms. The Meter Pulse frequency has been provisioned to the MTA.
- 2) The IPAT requests application of the meter pulse signal by the MTA:

```

RQNT 2367 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 7632
S: E/ps(lt=mpb, pd=150, pr=2 000, rep=25)
R: oc, hu, hf.

```

- 3) The MTA acknowledges the request.
- 4) The MTA begins applying meter pulses to the subscriber's access line.
- 5) In this example, the LE requested notification on operation completion in the original request to generate the fixed number of meter pulses. The MTA now notifies the IPAT that the operation is complete:

```

NTFY 12876 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
X: 7632
O: oc(E/ps(mpb)).

```

- 6) The IPAT acknowledges the event notification.
- 7) The IPAT relays the pulsed signal completion acknowledgement to the LE.




Figure B.9: Fixed meter pulse application, completed

B.7.4 Steady signal line treatment

B.7.4.1 Steady signal line treatment - Reverse polarity

This call flow illustrates a Steady Signal request in which the V5 LE has requested reverse polarity to be applied to the a-b POTS terminals.

- 1) The V5 LE includes a reverse polarity steady signal request in a message to the IPAT.
- 2) The IPAT converts the binary coded V5 message and maps the binary coded reverse polarity treatment message to the NCS It message and sends the line treatment message to the MTA:

```

RQNT 550 aaln/1@rgw.mso.net MGCP 1.0 NCS 1.0
S: E/ss(lt=rp).
  
```

- 3) The MTA acknowledges the signal request:

```

200 550 OK.
  
```

- 4) The MTA applies a reverse polarity to the a-b terminals for the aaln/1 line presence on the MTA.

Annex C (normative): Dynamic Quality of Service

In this annex, we provide additional detail on the usage of Dynamic Quality of Service (D-QoS) in NCS. We describe the expected MTA behaviour in more detail and include a state machine that the MTA may implement to support the D-QoS behaviour described. The IPcablecom Dynamic Quality of Service Specification (J.dqos) should be consulted for further details.

MTA, in implementing support for Dynamic Quality of Service needs to store and maintain D-QoS state on a per connection basis. Whenever D-QoS has been used for a connection, the endpoint will keep the following D-QoS information associated with the connection until it is deleted:

- **GateID:** The current GateID used for the connection.
- **ResourceID:** The current ResourceID used for the connection.
- **Last reservation:** The parameters for the most recent reservation for the connection. This includes classifiers as well as media parameters in both the send and receive direction.
- **Last commit:** The parameters for the most recent commit for the connection. This includes classifiers as well as media parameters in both the send and receive direction.
- **Reserve Destination:** An IP address and port that may be used to enable resource reservations where the remote address info is not yet known as explained below.
- **Gate Location** - The IP address and port where the D-QoS commit message should be sent to when using RSVP. The MTA learns this address through the RSVP QoS messages.

The GateID is the key to resource reservation. Once a GateID has been provided for a connection, a D-QoS state machine is created for the connection, and all of the above information will be maintained for the connection until it is either deleted, or a new GateID is provided. In the latter case, the D-QoS state machine and the above information is reset, and the old reservation is deleted (see note).

NOTE: If a ResourceID is included, and that ResourceID matches the old ResourceID, then the old reservation should not be deleted before the new one is made.

Resources can be reserved and committed independently in both the send and receive direction by the MTA. The send destination IP address and port as well as the source IP address are taken from the RemoteConnectionDescriptor, when a RemoteConnectionDescriptor has been provided. In that case, the MTA MUST use the following classifiers for the resource reservation and commit.

Table C.1

		MTA-o (J.112/RSVP)
Downstream/receive		
Source IP		IP(SDP-t)
Source Port		*
Destination IP		IP(SDP-o)
Destination Port		Port(SDP-o)
Upstream/send		
Source IP		IP(SDP-o)
Source Port		Port(o)
Destination IP		IP(SDP-t)
Destination Port		Port(SDP-t)

where:

- **IP(SDP-o)** refers to the media IP address in MTA-o's LocalConnectionDescriptor.
- **IP(SDP-t)** refers to the media IP address in MTA-o's RemoteConnectionDescriptor.

- **Port(SDP-o)** refers to the media port in MTA-o's LocalConnectionDescriptor.
- **Port(o)** refers to the source port MTA-o will be using when sending media on this connection. Note that this may or may not be the same as Port(SDP-o).

When a RemoteConnectionDescriptor has not yet been provided, the actual send destination IP address and port is unknown and the ReserveDestination address is therefore used instead. For the receive direction, the source IP address and port is wild-carded. This enables a reservation and a receive commit of the resource on the access link. The following classifiers **MUST** be used.

Table C.2

		MTA-o (J.112/RSVP)
Downstream/receive		
	Source IP	*
	Source Port	*
	Destination IP	IP(SDP-o)
	Destination Port	Port(SDP-o)
Upstream/send		
	Source IP	IP(SDP-o)
	Source Port	Port(o)
	Destination IP	IP(RD-o)
	Destination Port	Port(RD-o)

where:

- **IP(RD-o)** refers to the IP address in the ReserveDestination supplied.
- **IP(Port-o)** refers to the port number in the ReserveDestination supplied. If no port number is specified a default value of 9 applies.
- Once the actual send destination and receive source media addresses and port are known, the reservations are updated with the appropriate classifiers.
- When RSVP is used as the resource reservation protocol, the destination address used for the RSVP PATH message will be the ReserveDestination IP address supplied until a RemoteConnectionDescriptor is supplied.

NCS/D-QoS state machine

As explained above, the MTA maintains state for the Dynamic Quality of Service used on a connection. The state is derived from a state machine which is driven by the following:

- **Current state** which consists of the pair (SendQoSState, ReceiveQoSState), where each QoS state may be one of the following:
 - **N** - No resource reservation exists for the direction.
 - **R** - A resource reservation exists for the direction, but no resources are currently committed.
 - **C** - A resource reservation exists for the direction, some resources are currently committed.
 - **Connection mode** which is the NCS connection mode. The connection modes "Conference", "Network Loopback", and "Network Continuity Test" are not shown explicitly in the state machine, as they are all similar to "SendReceive". The connection mode "Replicate" is also not shown as it is similar to "SendOnly".
- **Resource Change** which is one or more of the following:
 - RemoteConnectionDescriptor IP address or port changes (classifier needs to be updated). This includes the case where it arrives for first time.
 - Codec changes.

- Ptime changes.
- etc.
- The **D-QoS** rules provided in clause 6.3.3.

As explained above, the state machine will be reinitialized when a new GateID is received. If a ResourceID is supplied as well and it is the same as the old ResourceID, the reservation(s) for the new state machine **MUST** be performed before the reservation(s) for the old state machine are released.

The set of possible *states* are:

- (N, N) Send resources not reserved, receive resources not reserved.
- (R, R) Send resources reserved, receive resources reserved.
- (C, R) Send resources reserved and committed, receive resources reserved.
- (R, C) Send resources reserved, receive resources reserved and committed.
- (C, C) Send resources reserved and committed, receive resources reserved and committed.
- (R, N) Send resources reserved, receive resources not reserved.
- (C, N) Send resources reserved and committed, receive resources not reserved.
- (N, R) Send resources not reserved, receive resources reserved.
- (N, C) Send resources not reserved, receive resources reserved and committed.

Once resources have been reserved and/or committed for a direction, a reservation for that direction will exist for the lifetime of the connection. The relationship between states and connection mode or D-QoS reservation parameters is shown in table C.3.

Table C.3

	SendState	RecvState
No Reserve/Commit parameter supplied - connection mode:		
inactive	R	R
sendonly, replcate	C	R
recvonly	R	C
sendrecv, confrnce, netwloop, netwttest	C	C
Reserve/Commit parameter supplied:		
sendresv	R	N, R (note)
recvresv	N, R (note)	R
snrcresv	R	R
sendcomt	C	N, R (note)
recvcomt	N, R (note)	C
snrccomt	C	C
NOTE: If resources have been reserved or committed previously for the direction, the state will be R, otherwise the state will be N.		

The actual state transition diagram is depicted in figure C.1.

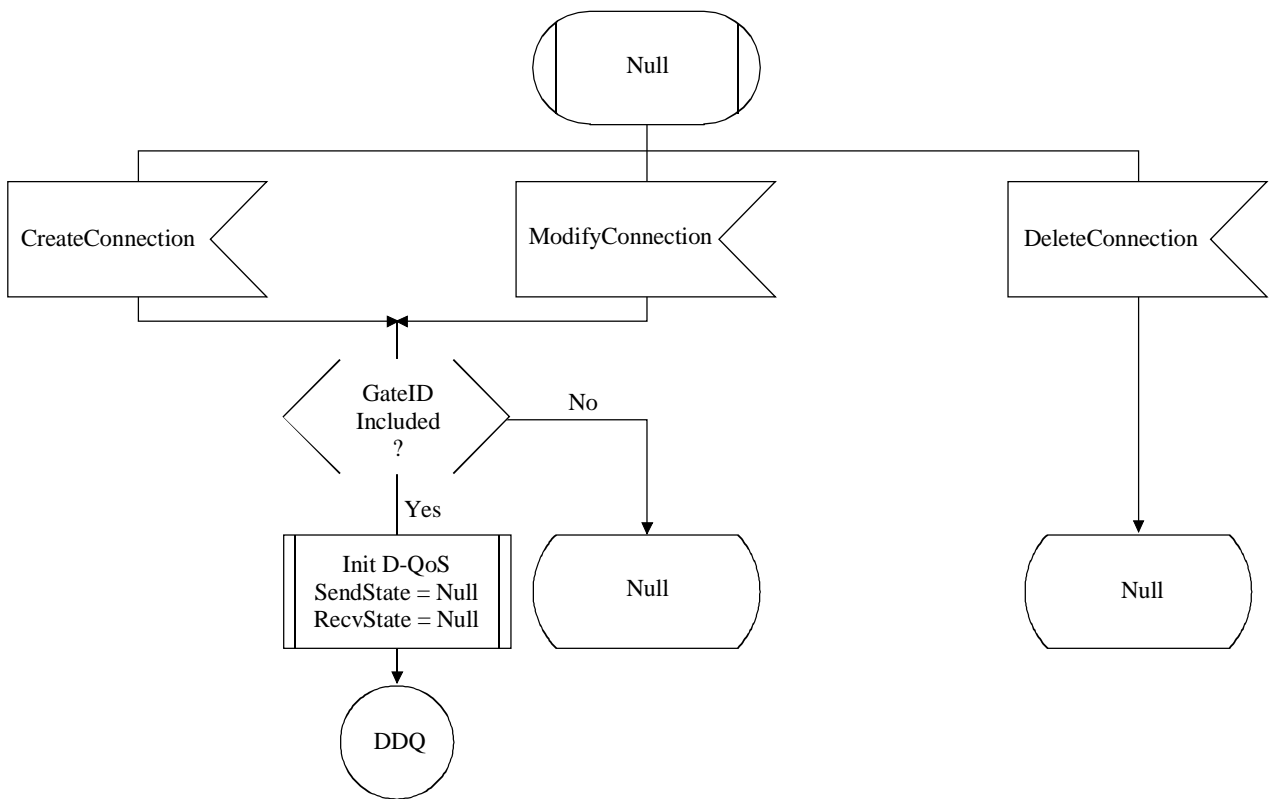


Figure C.1: NCS/D-QoS state diagram (1:2)




Figure C.2: NCS/D-QoS state diagram (2:2)

When executing the state machine, boolean variables will be set to indicate whether reserve, unreserve, commit, and uncommit operations are to be performed. The pseudo-code below then provides details on individual D-QoS procedures that are to be executed as indicated by these booleans. The following *actions* specify the D-QoS actions to be taken in each of these procedures:

- **SR** means a D-QoS Send Reservation will be performed.
- **RR** means a D-QoS Receive Reservation will be performed.
- **SC** means a D-QoS Send Commit will be performed.
- **RC** means a D-QoS Receive Commit will be performed.
- **SD** means a D-QoS Send Reservation Delete will be performed.
- **RD** means a D-QoS Receive Reservation Delete will be performed.
- **SU** means a D-QoS Send Uncommit, i.e. lower committed send resources to zero, will be performed.
- **RU** means a D-QoS Receive Uncommit, i.e. lower committed send resources to zero, will be performed.

SendReserve()

```
If <current resources reserved ≠ resources to reserve> then { -- skip reservation if existing
    reservation OK
        If <RemoteConnectionDescriptor provided> then
            SR(RemoteConnectionDescriptor) -- Use RemoteConnectionDescriptor classifier
        else if <ReserveDestination provided> then
            SR(ReserveDestination) -- Use ReserveDestination classifier, send to
                -- ReserveDestination if RSVP
        else ERROR
    }
```

ReceiveReserve()

```
If <current resources reserved ≠ resources to reserve> then {-- skip reservation if existing reservation
    OK
        If <RemoteConnectionDescriptor provided> then
            RR(RemoteConnectionDescriptor) -- Use RemoteConnectionDescriptor classifier
        else if <(J.112 QoS) or (RSVP and ReserveDestination provided )> then
            RR(*) -- Use wildcard classifier, send to ReserveDestination
                -- if RSVP
        else ERROR
    }
```

SendCommit()

```
If <current resources committed ≠ resources to commit> then {-- skip commit if existing OK
    If <RemoteConnectionDescriptor provided> then {
        If not <resources to commit ≤ resources reserved > then { -- old reservation does not
            SR(RemoteConnectionDescriptor) -- satisfy what is about to be
        } -- committed, so update reservation
        if <(J.112 QoS) or (RSVP and ReserveDestination provided )> then {
            SC(RemoteConnectionDescriptor) -- send to ReserveDestination if
                -- RSVP
        } else
        ERROR
    } else ERROR. -- Cannot commit send direction without RemoteConnectionDescriptor
}
```

ReceiveCommit()

```

If <current resources committed ≠ resources to commit> then {-- skip commit if existing OK
    If not <resources to commit ⊂ resources reserved> then {
        If <RemoteConnectionDescriptor provided> then
            RR(RemoteConnectionDescriptor)
        else if <(J.112 QoS) or (RSVP and ReserveDestination provided )> then
            RR(*) -- Use wildcard classifier, send to ReserveDestination if RSVP
        else
            ERROR
        }
        If <RemoteConnectionDescriptor provided> then
            RC(RemoteConnectionDescriptor)
        else if <(J.112 QoS) or (RSVP and ReserveDestination provided )> then
            RC(*) -- Use wildcard classifier, send to ReserveDestination if RSVP
        else
            ERROR
    }
}

```

SendReserveDelete()

```

If <send resources reserved> then
    SD() -- delete the reservation

```

ReceiveReserveDelete()

```

If <receive resources reserved> then
    RD() -- delete the reservation

```

SendUnCommit()

```

If <send resources commit> then
    SU() -- uncommit committed resources

```

ReceiveUnCommit()

```

If <receive resources committed> then
    RU() -- uncommit committed resources

```

State UpdateState(DoCommit, DoReserve, OldState)

```

If <DoCommit = true> then
    return Commit
else if <DoReserve = true> then
    return Reserve
else
    return OldState

```

Annex D (informative): Mode interactions

An MGCP connection can establish one or more media streams. These streams are either incoming (from a remote endpoint) or outgoing (generated at the handset microphone). The "connection mode" parameter establishes the direction and generation of these streams. When there is only one connection to an endpoint, the mapping of these streams is straightforward; the handset plays the incoming stream over the handset speaker and generates the outgoing stream from the handset microphone signal, depending on the mode parameter.

However, when several connections are established to an endpoint, there can be many incoming and outgoing streams. Depending on the connection mode used, these streams may interact differently with each other and the streams going to/from the handset.

Table D.1 describes how different connections should be mixed when one or more connections are concurrently "active". An active connection is here defined as a connection that is in one of the following modes:

- "send/receive";
- "send only";
- "receive only";
- "replicate";
- "conference".

Connections in "network loopback", "network continuity test", or "inactive" modes are not affected by connections in the "active" modes. The table uses the following conventions:


- A_{in} is the incoming media stream from Connection A.
- B_{in} is the incoming media stream from Connection B.
- H_{in} is the incoming media stream from the Handset Microphone.
- A_{out} is the outgoing media stream to Connection A.
- B_{out} is the outgoing media stream to Connection B.
- H_{out} is the outgoing media stream to the Handset earpiece.
- NA indicates No Stream whatever.

Table D.1

		Connection A Mode						
		sendonly	Recvonly	sendrecv	confrnce	inactive	netwloop/ netwtest	replicate
Connection B Mode	sendonly	A _{out} = H _{in} B _{out} = H _{in} H _{out} =NA	A _{out} =NA B _{out} =H _{in} H _{out} = A _{in}	A _{out} =H _{in} B _{out} =H _{in} H _{out} =A _{in}	A _{out} =H _{in} B _{out} =H _{in} H _{out} =A _{in}	A _{out} = NA B _{out} = H _{in} H _{out} =NA	A _{out} =A _{in} B _{out} = H _{in} H _{out} =NA	A _{out} = H _{in} B _{out} = H _{in} H _{out} =NA
	recvonly		A _{out} = NA B _{out} = NA H _{out} =A _{in} +B _{in}	A _{out} = H _{in} B _{out} = NA H _{out} =A _{in} +B _{in}	A _{out} =H _{in} B _{out} =NA H _{out} =A _{in} +B _{in}	A _{out} =NA B _{out} =NA H _{out} = B _{in}	A _{out} = A _{in} B _{out} = NA H _{out} = B _{in}	A _{out} = H _{in} +B _{in} B _{out} = NA H _{out} = B _{in}
	sendrecv			A _{out} = H _{in} B _{out} = H _{in} H _{out} =A _{in} +B _{in}	A _{out} = H _{in} B _{out} =H _{in} H _{out} =A _{in} +B _{in}	A _{out} = NA B _{out} = H _{in} H _{out} = B _{in}	A _{out} = A _{in} B _{out} = H _{in} H _{out} = B _{in}	A _{out} = H _{in} +B _{in} B _{out} = H _{in} H _{out} = B _{in}
	confrnce				A _{out} =H _{in} +B _{in} B _{out} =H _{in} +A _{in} H _{out} =A _{in} +B _{in}	A _{out} = NA B _{out} = H _{in} H _{out} = B _{in}	A _{out} = A _{in} B _{out} = H _{in} H _{out} = B _{in}	A _{out} = H _{in} +B _{in} B _{out} = H _{in} H _{out} = B _{in}
	inactive					A _{out} = NA B _{out} = NA H _{out} = NA	A _{out} = A _{in} B _{out} = NA H _{out} =NA	A _{out} = H _{in} B _{out} = NA H _{out} =NA
	netwloop/ netwtest netwtest						A _{out} = A _{in} B _{out} = B _{in} H _{out} =NA	A _{out} = H _{in} B _{out} = B _{in} H _{out} =NA
	replicate							A _{out} = H _{in} B _{out} = H _{in} H _{out} =NA


If there are three or more "active" channels they will still interact as defined in table D.1 with the outgoing media streams mixed for each interaction. (Union of all streams) If internal resources are used up and the streams cannot be mixed, the gateway should return a resources Not available error.

These connections can be graphically represented as such:




For example, if Connection A is Sendrecv, Connection B is confrnce, and Connection C is recvonly, from the above table the outputs in each mode will be:


A to B Interaction: B_{out}= H_{in} A_{out}=H_{in} H_{out}=A_{in}+B_{in}



A to C Interaction: $A_{out} = H_{in}$ $C_{out} = NA$ $H_{out} = A_{in} + C_{in}$




B to C Interaction: $B_{out} = H_{in}$ $C_{out} = NA$ $H_{out} = B_{in} + C_{in}$



Taking the Union of all streams in each output we get:

- $A_{out} = H_{in}$
- $B_{out} = H_{in}$
- $C_{out} = NA$
- $H_{out} = B_{in} + A_{in} + C_{in}$.



For clarity, the table described above is repeated below in graphical form (excluding the "replicate" mode):

Table D.2

		Connection A Mode (top)					
		sendonly	recvonly	sendrecv	confrnce	inactive	netwloop/ netwtest
Connection B Mode (Left)	sendonly						
	recvonly						
	sendrecv						
	confrnce						
	inactive						
	netwloop/ netwtest						

Annex E (informative): Compatibility information

This annex provides NCS protocol compatibility information.

MGCP Compatibility

NCS is a profile of MGCP 1.0, however NCS has introduced a couple of additions as well. The following lists NCS additions that are currently not included in MGCP:

- **Endpoint Naming Scheme** - The rules for wildcarding are more restrictive than in MGCP.
- **Embedded ModifyConnection** - A new Embedded ModifyConnection action has been introduced.
- **Dynamic Quality of Service** - IPCablecom Security services are supported in NCS. This affects the LocalConnectionOptions, Capabilities, and SDP. Also, a new return parameter; ResourceID, is added for CreateConnection and ModifyConnection.
- **Security** - IPCablecom Security services are supported in NCS. This affects the LocalConnectionOptions, Capabilities, and SDP.
- **Endpoint Name Retrieval** - The AuditEndpoint command has been extended with a capability to return the number of endpoints that match a wildcard as well as mechanism for block-wise retrieval of these endpoint names. Besides extending the AuditEndpoint command, this implies the introduction of two new parameter names; MaxEndPointIds, and NumEndPoints.
- **Supported Versions** - The RestartInProgress response and the AuditEndpoint command have been extended with a VersionSupported parameter to enable Call Agents and gateways to determine which protocol versions each support.
- **Error Codes** - Two new error codes have been introduced; 532 and 533.
- **Usage of SDP** - A new SDP usage profile is included in NCS. Most notably, the profile and all example use specifically require strict SDP compliance, regardless of the usefulness of the included fields. Also, IPCablecom specific extensions have been added to SDP.
- **Provisional Response** - Additional detail and specification of the provisional response mechanism has been included in NCS. A Response Acknowledgement response (000) has been introduced, an empty ResponseAck parameter has been permitted in final responses that follow provisional responses, and a procedure for the mechanism specified.
- **Signal Parameters** - Signal parameter syntax has been extended to allow for the usage of balanced parenthesis within signal parameters. All Time-Out signals can have their time-out value altered by a signal parameter.
- **Event Packages** - NCS introduces a set of new event packages.

Finally, it should be noted, that NCS provides interpretations of and in some cases additional specification or clarification of the base MGCP protocol behaviour that may or may not reflect the intended MGCP behaviour.

Annex F (informative): Metering support for IPCablecom NCS

F.1 Objectives

As described in EPC-RequDoc-V10-220501 [19], hardware metering is a requirement in the support of analogue lines in an IP Cable environment. This annex describes a package for the automatic transmission of hardware metering pulses on analogue lines. It also includes metering-specific call flows.

NOTE: The description of a standalone automatic meter package given by this annex and that described by annex B are the same intentionally, and should remain aligned. The equivalence of the meter pulse signals as described in annex B and that described in this annex is directly mapped; E/ps(lt=em) maps directly to am/em and E/ps(mpb) maps directly to am/mpb respectively. These signals accept the same parameter usage in both packages.

A consideration in generating this package was the decoupling of the Media Gateway from currency knowledge. The charge unit varies per market. The Media Gateway should not have to know the value of a pulse (currency units).

F.2 Automatic metering package

The Automatic Metering Package is designed to meet the requirements of Media Gateways with analogue lines configured for general purpose telephony, adding a capability for automatic transmission of metering pulses.

Pulse characteristics (pulse type, pulse duration, minimum pause duration) are market-dependent (see EN 300 001 [18]) and do not change during a call. By not including pulse characteristics in the MGCP message the package retains the capability to support any type of metering pulse in any market. This package assumes that pulse characteristics are provisioned (MIB) on the Media Gateway.

This package assumes that accumulation is the task of the CPE device. This package does not require the Media Gateway to keep track of the number of pulses generated.

This package assumes gateways are reliable with respect to the generation of pulses. It does not include feedback (events, properties, statistics) on the number of pulses actually generated during a call.

F.2.1 Package name

Package Name: am.

Version: 1.

Metering signals and events MUST always be prefixed with the "am" package name.

F.2.2 Local connection options

None.

F.2.3 Events and signals

This package introduces two signals.

Table F.1: Signals in the metering package

Symbol	Definition	R	Type	Duration
em	enable metering		OO	n/a
mpb	meter pulse burst		BR	n/a

Legend:

R: an "x" appears in this column if the event can be requested by the Call Agent. Alternatively, an "S" may be included if the event-state may be audited. A "C" indicates that the event can be detected on a connection.

Type: if nothing appears in this column for an event, then the event cannot be signalled on command by the Call Agent. Otherwise, the following symbols identify the type of event:

- OO On/Off signal.
- TO Timeout signal.
- BR Brief signal.

Duration: specifies the duration of TO signals. If a duration is left unspecified then the default timeout will be assumed to infinite.

F.2.3.1 Meter pulse burst signal

Signal Name: am/mpb.

Signal Type: Brief.

The meter pulse signal is used to signal call attempt, call set-up and add-on charges. It requests the generation of a fixed number of metering pulses on the analogue line. Note that the meter pulse signal may also be used to request the generation of a single meter pulse.

Additional parameters:

- pulse count:
 - ParameterID: rep;
 - Type: integer, rep > 0;
 - Default value: 1.

This parameter specifies the number of metering pulses to be applied on the line. The MTA SHALL generate pulses until the pulse count is reached.

The default value of this parameter, which SHALL apply if the parameter is omitted, is 1.

- Pulse repetition interval:
 - ParameterID: pr;
 - Type: integer, pr > 0;
 - Default: 1 000.

This parameter specifies the interval between repetitions of metering pulses on the line, in ms. It represents the time that SHOULD elapse between the leading edge of a pulse and the leading edge of the succeeding pulse.

The default value of this parameter, which SHALL apply if the parameter is omitted, is 1 000 ms.

A meter pulse burst signal request may be included in a signal request that enables metering pulse generation, for example, to apply an initial charge to a call. When this occurs, the MTA SHALL apply the metering pulse burst to the endpoint completely, and then start generating normal metering pulses.

Because the metering pulse burst signal is a brief signal type, all pulses specified to for the request (**rep=n**) SHALL be applied, even if the subscriber goes on-hook during the burst.

It is considered an error when a MTA receives a meter pulse burst signal and the set is on hook. When such attempts are made an error code 402 (phone on hook) SHALL be returned.

The am/mpb signal SHALL be applied to endpoints, NOT to connections.

F.2.3.2 Enable metering signal

Signal Name: am/em.

Signal Type: On/Off.

This signal starts the automatic generation of metering pulses on the analogue line. It is used to signal a regular, time-based call charge. The first pulse of a call charge SHALL be pulsed out immediately after the em signal is received.

Additional parameters:

- Pulse repetition interval:
 - ParameterID: pr;
 - Type: integer, pr > 0;
 - Default: 1 000.

This parameter specifies the interval between repetitions of metering pulses on the line, in ms. It represents the time that SHOULD elapse between the leading edge of a pulse and the leading edge of the succeeding pulse. The MTA SHALL continue to generate pulses until it receives a new am/em signal or the em signal is explicitly turned off. If a line goes on hook, the MTA SHOULD disable meter pulses in anticipation of a new call set-up (CPE going back off hook for a new call).

The default value of this parameter, which SHALL apply if the parameter is omitted, is 1 000 ms.

Enable metering signals are mutually exclusive; only one enable metering signal SHALL be active at a time. If a new am/em signal arrives it SHALL supersede (replace) any previous am/em signal.

A metering pulse burst signal request may occur during a call in progress, for example to take account of a chargeable subscriber action. When this occurs, the MTA SHALL suspend normal metering pulse generation, and apply the metering pulse burst signals. The MTA then shall resume normal metering pulse generation without requiring a new "enable metering" request from the Call Agent. The Call Agent MUST account for any normal metering pulses missed during the burst by including the missed pulses in the burst count.

It is considered an error when a MTA receives an enable metering signal and the set is on hook. When such attempts are made an error code 402 (phone on hook) SHALL be returned.

The syntax to turn the enable metering signal off is am/em(-). When an enable metering off signal is received in on hook state no error SHALL be returned.

The am/em signal SHALL be applied to endpoints, NOT to connections.

F.2.4 Properties

None.

F.2.5 Statistics

None.

F.2.6 Procedures

None.

F.3 Use cases, Example call flows

F.3.1 Meter pulse while off hook

The Call Agent instructs the MTA to apply a single pulse. If omitted, the "rep" parameter defaults to 1. The set is off hook.

```
RQNT 309 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 860
S: am/mpb
```

The MTA confirms.

```
200 309 ok
```

F.3.2 Meter pulse while on hook

The Call Agent instructs the MTA to apply a single pulse while the set is on hook.

```
RQNT 310 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 870
S: am/mpb
```

The MTA rejects the request.

```
402 310 phone on hook
```

F.3.3 Regular call charge

The Call Agent instructs the MTA to apply a regular call charge of one pulse every 12 s.

```
RQNT 311 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 880
S: am/em(pr=12000)
```

The MTA confirms.

```
200 311 ok
```

F.3.4 Call set-up charge

The Call Agent instructs the MTA to apply a pulse burst of 33 pulses.

```
RQNT 321 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 881
S: am/mpb(rep=33)
```

The MTA confirms.

```
200 321 ok
```

Afterwards the Call Agent instructs the MTA to apply a regular call charge of one pulse every 5 s.

```
RQNT 322 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 882
S: am/em(pr=5000)
```

The MTA confirms.

```
200 322 ok
```


Note that the Call Agent has the option to apply both signals in one request:

```
RQNT 323 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 883
S: am/mpb(rep=33), am/em(pr=5000)
```

The MTA confirms.

```
200 323 ok
```

F.3.5 Mid-call tariff change

The Call Agent instructs the MTA to apply a regular call charge of one pulse every 8 s.

```
RQNT 331 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 884
S: am/em(pr=8000)
```

The MTA confirms.

```
200 331 ok
```

Later, as the call progresses into a different time of day, the tariff changes. The Call Agent instructs the MTA to apply a pulse every 12 s.

```
RQNT 332 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 885
S: am/em(pr=12000)
```

The MTA confirms.

```
200 332 ok
```

F.3.6 Mid-call add-on charge

Say the call is initially routed to an announcement. The Call Agent instructs the MTA to apply a pulse every 10 s.

```
RQNT 341 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 886
S: am/em(pr=10000)
```

The MTA confirms.

```
200 341 ok
```

Later, as the call is transferred to an operator, an add-on charge is applied. The Call Agent instructs the MTA to apply a one-time pulse burst of 20 pulses, without affecting the regular call charge.

```
RQNT 342 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
X: 887
S: am/mpb(rep=20)
```

The MTA confirms.

```
200 342 ok
```

F.3.7 End of call

At the end of the call, the Call Agent instructs the MTA to delete the connection and turn the regular call charge off.

```
DLCX 351 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
C: abcd
S: am/em(-)
```

The MTA confirms.

```
250 351 ok
```

F.3.8 Audit endpoint

Brief signals do not have an auditable state. Per MGCP specification currently playing Brief signals are not included in the response to a signal request audit.

The state of On/Off signals is an auditable property. If the Audit Endpoint command asks for RequestedInfo=SignalRequests, the MTA MUST return a list of the On/Off signals that are currently "On" for the endpoint (with or without parameters).

The Call Agent audits the endpoint.

```
AUEP 361 aaln/1@mg23.whatever.net MGCP 1.0 NCS 1.0
F: S
```

The response indicates that a regular call charge signal is on.

```
200 361 ok
S: am/em(pr=10000)
```

F.4 Terms

Charge: number of charge units (for the usage of a chargeable event (telecommunication service)).

Charge unit: base element for the charging process, expressed as meter-pulse units or as currency value.


Add-on charge: single additional charge which does not change the current tariff.

Tariff: set of parameters used for charging purposes to calculate the numbering charge units for the telecommunication service or a group of telecommunication services used. A tariff consists of a tariff sequence.

Tariff sequence: list of up to 4 consecutive subtariffs which has to be applied for the charging of the communication event. The subtariffs are applied at the start of the communication event and are applied consecutively according to the list of the subtariffs. The last subtariff may have an unlimited duration.

Subtariff: within a tariff sequence, a charge unit per time unit. Each subtariff has an individual duration and an individual charge unit.

Meter pulse: a periodic, cadenced signal with one on- and one off-period. The three most common types of metering pulses are: 12 kHz pulse, 16 kHz pulse and reverse polarity pulse.



MIB: Management Information Base.

Pulse repetition interval: varies with the charge; the higher the charge, the shorter the pulse repetition interval.

On-period (pulse): of fixed length; its duration depends, however, on national specifications. See EN 300 001 [18], clause 1.7.8.

Off-period (pause): varies with the pulse repetition interval; its minimum duration depends on national specifications. See EN 300 001 [18], clause 1.7.8.

Annex G (informative): Bibliography

ETSI TS 101 909-2: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 2: Architectural framework for the delivery of time critical services over cable Television networks using cable modems".

ETSI TS 101 909-5: "Access and Terminals (AT); Digital Broadband Cable Access to the Public Telecommunications Network; IP Multimedia Time Critical Services; Part 5: Dynamic Quality of Service for the Provision of Real Time Services over Cable Television Networks using Cable Modems".

ITU-T Recommendation H.225: "Call signalling protocols and media stream packetization for packet-based multimedia communication systems".

ITU-T Recommendation H.245: "Control protocol for multimedia communication".

ITU-T Recommendation H.323: "Framework and wire-protocol for multiplexed call signalling transport".

ITU-T Recommendation T.30: "Procedures for document facsimile transmission in the general switched telephone network".

ITU-T Recommendation V.8: "Procedures for starting sessions of data transmission over the public switched telephone network".

ITU-T Recommendation V.21: "300 bits per second duplex modem standardized for use in the general switched telephone network".

ETSI TR 101 963: "Access and Terminals (AT); Report on Requirements of European Cable Industry for implementation of Europacket cable/IPCablecom technologies; Identification of high level requirements and establishment of priorities".

Arango, Dugan, Elliot, Huitema, Pickett, Foster, Andreasen, Basic MGCP Packages, Work in Progress, draft-foster-mgcp-basic-packages-00.txt, Internet Engineering Task Force, January 2001.

ETSI ES 201 296: "Integrated Services Digital Network (ISDN); Signalling System No.7; ISDN User Part (ISUP); Signalling aspects of charging".

ETSI TS 101 868: "Services and Protocols for Advanced Networks (SPAN); Recommended V5 PSTN mapping".

ITU-T Recommendation V.18: "Operational and interworking requirements for DCEs operating in the text telephone mode".

List of ITU-T Recommendations referring to IP Cablecom:

ITU-T Recommendation J.160: "Architectural framework for the delivery of time critical services over cable television networks using cable modems".

ITU-T Recommendation J.161: "Audio codec requirements for the provision of bidirectional audio service over cable television networks using cable modems".

ITU-T Recommendation J.162: "Network call signalling (NCS) MIB requirements".

ITU-T Recommendation J.163: "Dynamic quality of service for the provision of real time services over cable television networks using cable modems".

ITU-T Recommendation J.164: "Event Message requirements for the support of real-time services over cable television networks using cable modems".

ITU-T Recommendation J.165: "IPCablecom Internet Signalling Transport Protocol".

ITU-T Recommendation J.166: "IPCablecom Management information base (MIB) framework".

ITU-T Recommendation J.167: "Media terminal adapter (MTA) device provisioning requirements for the delivery of real-time services over cable television networks using cable modems".

ITU-T Recommendation J.168: "IPcablecom media terminal adapter (MTA) MIB requirements".

ITU-T Recommendation J.169: "IPcablecom network call signalling (NCS) MIB requirements".

ITU-T Recommendation J.170: "IPcablecom Security specification".

ITU-T Recommendation J.171: "IPcablecom Trunking Gateway Control Protocol (TGCP)".

History

Document history		
V1.1.1	July 2001	Publication
V1.2.1	February 2002	Publication
V1.2.2	July 2002	Publication
V1.3.1	December 2002	Publication