

PacketCable™ 1.5 Specifications

Network-Based Call Signaling Protocol

PKT-SP-NCS1.5-C01-191120

CLOSED

Notice

This PacketCable™ specification is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs® in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third-party documents, including open source licenses, if any.

© 2004-2019 Cable Television Laboratories, Inc.
All rights reserved

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs and its members shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	PKT-SP-NCS1.5-C01-191120			
Document Title:	Network-Based Call Signaling Protocol			
Revision History:	I01 - Issued January 28, 2005 I02 - Issued August 12, 2005 I03 - Issued April 12, 2007 I04 - Issued April 12, 2012 C01 - Released November 20, 2019			
Date:	November 20, 2019			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author Only	CL/Member	CL/ PacketCable/ Vendor	Public

Key to Document Status Codes:

Work in Progress	An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
Draft	A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
Issued	A generally public document that has undergone Member and Technology Supplier review, cross-vendor interoperability, and is for Certification testing if applicable. Issued Specifications are subject to the Engineering Change Process.
Closed	A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks:

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

Contents

1	STATUS OF THIS DOCUMENT	1
1.1	SPECIFICATION LANGUAGE	1
2	REFERENCES.....	2
2.1	NORMATIVE REFERENCES	2
2.2	INFORMATIVE REFERENCES	3
2.3	REFERENCE ACQUISITION	3
3	TERMS AND DEFINITIONS.....	4
4	ABBREVIATIONS AND ACRONYMS	8
5	SCOPE	15
6	INTRODUCTION.....	17
6.1	RELATION WITH H.323 STANDARDS.....	18
6.2	RELATION WITH IETF STANDARDS	18
6.3	RELATION TO RFC 3435 AND ABNF GRAMMAR	19
7	MEDIA GATEWAY CONTROL INTERFACE (MGCI).....	20
7.1	MODEL AND NAMING CONVENTIONS.....	20
7.1.1	<i>Endpoint Names.....</i>	20
7.1.2	<i>Call Names.....</i>	22
7.1.3	<i>Connection Names</i>	22
7.1.4	<i>Names of Call Agents and Other Entities</i>	22
7.1.5	<i>Digit Maps</i>	23
7.1.6	<i>Events and Signals.....</i>	25
7.2	SDP USE	27
7.3	GATEWAY CONTROL FUNCTIONS.....	27
7.3.1	<i>NotificationRequest.....</i>	29
7.3.2	<i>Notifications.....</i>	35
7.3.3	<i>CreateConnection</i>	36
7.3.4	<i>ModifyConnection.....</i>	42
7.3.5	<i>DeleteConnection (From the Call Agent).....</i>	44
7.3.6	<i>DeleteConnection (From the Embedded Client).....</i>	46
7.3.7	<i>DeleteConnection (Multiple Connections From the Call Agent).....</i>	46
7.3.8	<i>Auditing</i>	47
7.3.9	<i>Restart in Progress</i>	51
7.4	STATES, FAILOVER AND RACE CONDITIONS	52
7.4.1	<i>Recaps and Highlights</i>	52
7.4.2	<i>Retransmission, and Detection of Lost Associations</i>	53
7.4.3	<i>Race Conditions.....</i>	56
7.5	RETURN CODES AND ERROR CODES	66
7.6	REASON CODES.....	67
7.7	USE OF LOCAL CONNECTION OPTIONS AND CONNECTION DESCRIPTORS	67
7.7.1	<i>RFC 2833 Negotiation.....</i>	70
7.7.2	<i>Selecting Codec Packetization Interval for T.38.....</i>	70
7.7.3	<i>Remote IP and Port Negotiation.....</i>	71
7.7.4	<i>LCO Use with FXR Package.....</i>	71
7.7.5	<i>LCO with FXR and VBD Being Used by CMS at Same Time</i>	71
7.7.6	<i>LCO format for V.152 and RFC 2198 Redundancy use.....</i>	72
8	MEDIA GATEWAY CONTROL PROTOCOL	73
8.1	GENERAL DESCRIPTION	73

8.2	COMMAND HEADER	73
8.2.1	<i>Command Line</i>	74
8.2.2	<i>Parameter Lines</i>	75
8.3	RESPONSE HEADER FORMATS	86
8.3.1	<i>CreateConnection</i>	88
8.3.2	<i>ModifyConnection</i>	88
8.3.3	<i>DeleteConnection</i>	89
8.3.4	<i>NotificationRequest</i>	89
8.3.5	<i>Notify</i>	89
8.3.6	<i>AuditEndpoint</i>	89
8.3.7	<i>AuditConnection</i>	89
8.3.8	<i>RestartInProgress</i>	90
8.4	SESSION DESCRIPTION ENCODING	90
8.4.1	<i>Use of SDP</i>	90
8.4.2	<i>SDP Parameters Common to both Audio (m=audio media line) and Image (m=image media line) Service Use</i>	91
8.4.3	<i>SDP Audio Service Use</i>	94
8.4.4	<i>SDP Image Service Use for T.38</i>	98
8.5	TRANSMISSION OVER UDP	101
8.5.1	<i>Reliable Message Delivery</i>	101
8.5.2	<i>Retransmission Strategy</i>	101
8.5.3	<i>Maximum Datagram Size, Fragmentation and Reassembly</i>	102
8.6	PIGGY-BACKING	102
8.7	TRANSACTION IDENTIFIERS AND THREE WAYS HANDSHAKE	103
8.8	PROVISIONAL RESPONSES	104
9	SECURITY	106
APPENDIX A.	EVENT PACKAGES	107
A.1	ANALOG ACCESS LINES	107
A.2	LINE PACKAGE	107
A.3	FAX PACKAGE	113
A.4	VOIP METRICS PACKAGE	145
A.5	VBD AND GPMD PACKAGES	158
APPENDIX B.	CONNECTION MODE	191
APPENDIX C.	DYNAMIC QUALITY-OF-SERVICE	195
C.1	INTRODUCTION	195
C.2	NCS/D-QoS STATE MACHINE	196
APPENDIX D.	EXAMPLE COMMAND ENCODINGS	202
D.1	NOTIFICATIONREQUEST	202
D.2	NOTIFY	202
D.3	CREATECONNECTION	202
D.4	MODIFYCONNECTION	204
D.5	DELETECONNECTION (FROM THE CALL AGENT)	205
D.6	DELETECONNECTION (FROM THE EMBEDDED CLIENT)	205
D.7	DELETECONNECTION (MULTIPLE CONNECTIONS FROM THE CALL AGENT)	205
D.8	AUDITENDPOINT	205
D.9	AUDITCONNECTION	206
D.10	RESTARTINPROGRESS	207
APPENDIX E.	EXAMPLE CALL FLOW	209
APPENDIX F.	COMPATIBILITY INFORMATION	216
F.1	MGCP COMPATIBILITY	216

APPENDIX G.	ABNF GRAMMAR FOR NCS	217
APPENDIX H.	ACKNOWLEDGEMENTS	226
APPENDIX I.	REVISION HISTORY	227

List of Figures

FIGURE 1 - RELATION TO H.323 STANDARDS	18
FIGURE 2 - RELATION TO IETF STANDARDS	19
FIGURE 3 - DIALED DIGITS - EXAMPLES	23
FIGURE 4 - RETRANSMISSION ALGORITHM	55
FIGURE 5 - QUARANTINE LIST PROCEDURES	59
FIGURE 6 - CONTROL OF MEDIA STREAMS BY CONNECTION MODE - GRAPHICAL REPRESENTATION.....	193
FIGURE 7 - GRAPHICAL REPRESENTATION OF CONNECTIONS	194
FIGURE 8 - INDIVIDUAL MEDIA STREAMS ON CONNECTIONS	194
FIGURE 9 - A TO B TO C INTERACTION	194
FIGURE 10 - NCS/D-QoS STATE DIAGRAM (1:2).....	198
FIGURE 11 - NCS/D-QoS STATE DIAGRAM (2:2).....	199

List of Tables

TABLE 1 - PACKAGES FOR EMBEDDED CLIENT END-POINT TYPES	25
TABLE 2 - LEGAL COMBINATIONS OF EVENTS AND ACTIONS	32
TABLE 3 - DEFAULT RESOURCE RESERVATION VALUES	39
TABLE 4 - RETURN CODE DEFINITIONS	66
TABLE 5 - REASON CODE DEFINITIONS	67
TABLE 6 - REQUESTED VERB CODINGS	74
TABLE 7 - EXAMPLE NAME CODING.....	75
TABLE 8 - PARAMETER DEFINITIONS.....	76
TABLE 9 - ASSOCIATION OF PARAMETERS WITH COMMANDS.....	77
TABLE 10 - DQoS RESOURCE RESERVATION PARAMETER VALUES.....	79
TABLE 11 - CONNECTION PARAMETERS	81
TABLE 12 - CONNECTION MODE.....	82
TABLE 13 - EVENT NAME EXAMPLES.....	82
TABLE 14 - EVENT RANGE AND WILDCARD NOTATION	82
TABLE 15 - "ALL" AND "CURRENT" CONNECTION NOTATION.....	82
TABLE 16 - REQUESTED EVENTS ACTIONS	83
TABLE 17 - REQUESTEDINFO PARAMETER VALUES.....	85
TABLE 18 - ASSOCIATION OF RESPONSE HEADER PARAMETERS AND COMMANDS.....	87
TABLE 19 - LINE PACKAGE CODES FOR EVENTS AND SIGNALS	108
TABLE 20 - CONTROL OF MEDIA STREAMS BY CONNECTION MODE	192
TABLE 21 - CLASSIFIERS FOR RESOURCE RESERVATION AND COMMIT: REMOTE CONNECTION DESCRIPTOR PROVIDED	195
TABLE 22 - CLASSIFIERS FOR RESOURCE RESERVATION AND COMMIT: REMOTE CONNECTION DESCRIPTOR NOT PROVIDED	196
TABLE 23 - STATES RELATED TO CONNECTION MODES AND DQoS RESERVATION PARAMETERS	197
TABLE 24 - EXAMPLE CALL FLOW	209

1 STATUS OF THIS DOCUMENT

This document is considered part of the PacketCable™ specification. The document is based on MGCP 1.0 [1], which is an IETF Informational RFC.

1.1 Specification Language

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

"MUST"	This word or the adjective "REQUIRED" means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood, and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood, and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

The legal/regulatory classification of IP-based voice communications provided over cable networks and otherwise, and the legal/regulatory obligations, if any, borne by providers of such voice communications, are not yet fully defined by appropriate legal and regulatory authorities. Nothing in this specification is addressed to, or intended to affect, those issues. In particular, while this document uses standard terms such as "call," "call signaling," "telephony," etc., it will be evident from this document that while a PacketCable network performs activities analogous to these PSTN functions, the manner by which it does so differs considerably from the manner in which they are performed in the PSTN by telecommunications carriers. These differences may be significant for legal/regulatory purposes.

2 REFERENCES

2.1 Normative References

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

- [1] IETF RFC 3435, Media Gateway Control Protocol (MGCP) Version 1.0, January 2003.
- [2] IETF RFC 1889, RTP: A Transport Protocol for Real-Time Applications, January 1996.
- [3] IETF RFC 3551, RTP Profile for Audio and Video Conferences with Minimal Control, July 2003.
- [4] IETF RFC 2327, Session Description Protocol (SDP), April 1998.
- [5] IETF RFC 2974, SAP - Session Announcement Protocol, October 2000.
- [6] IETF RFC 2543, SIP: Session Initiation Protocol, March 1999.
- [7] IETF RFC 2326, Real-time Streaming Protocol (RTSP), April 1998.
- [8] TCP/IP Illustrated, Volume 1, *The Protocols*, Stevens, W. Richard, Addison-Wesley, 1994.
- [9] RTP Parameters, <http://www.Iana.org/assignments/rtp-parameters/>.
- [10] Bellcore, Bellcore Notes on the Networks, SR-2275.
- [11] PacketCable 1.5 Audio/Video Codecs Specification, PKT-SP-CODEC1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
- [12] PacketCable 1.5 Dynamic Quality of Service Specification, PKT-SP-DQOS1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
- [13] PacketCable 1.5 Security Specification, PKT-SP-SEC1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
- [14] IETF RFC 2821/STD0010, J. Klensin, Editor, Simple Mail Transfer Protocol, April 2001.
- [15] IETF RFC 1122/STD0003, Requirements for Internet Hosts -- Communication Layers, October 1989.
- [16] IETF STD0013, Domain Names – Implementations and Specifications, November 1987.
- [17] IETF RFC 2231, MIME Parameter Value and Encoded Word Extensions: Character Sets, November 1997.
- [18] IETF RFC 2234, Augmented BNF for Syntax Specifications: ABNF, November 1997.
- [19] IETF RFC 2833, RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals, May 2000.
- [20] GR-506, LSSGR: Signaling for Analog Interfaces, November 1996.
- [21] ITU-T Recommendation T.30, Procedures for Document Facsimile Transmission in the General Switched Telephone Network, April 1999.
- [22] ITU-T Recommendation V.8, Procedures for Starting Sessions of Data Transmission over the Public Switched Telephone Network, November 2000.
- [23] ITU-T Recommendation V.21, 300 Bits Per Second Duplex Modem Standardized for Use in the General Switched Telephone Network, November 1988.
- [24] 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, October 1996.
- [25] ITU-T Recommendation V.18, Operational and interworking requirements for DCEs operating in the text telephone mode, November 2000.

- [26] IETF RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998.
- [27] ITU-T Recommendation T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", March 2002.
- [28] IETF RFC 3611, "RTP Control Protocol Extended Reports (RTCP XR)", November 2003.
- [29] IETF RFC 2198, RTP Payload for Redundant Audio Data, September 1997.
- [30] ITU-T V.152, Procedures for supporting voice-band data over IP networks, January 2005.

2.2 Informative References

- [31] ITU-T Recommendation Q.761, Functional Description of the ISDN User Part of Signalling System No. 7, (Malaga-Torremolinos, 1984; modified at Helsinki, 1993).
- [32] ITU-T Recommendation Q.762, General Function of messages and Signals of the ISDN User part of Signalling System No. 7, (Malaga-Torremolinos, 1984; modified at Helsinki, 1993).
- [33] ITU-T Recommendation H.323, Visual Telephone Systems and Equipment for Local Area Networks Which Provide a Non-guaranteed Quality-of-Service, November 2000.
- [34] ITU-T Recommendation H.225, Call Signaling Protocols and Media Stream Packetization for Packet Based Multimedia Communications Systems, November 2000.
- [35] ITU-T Recommendation H.245, Line Transmission of Non-telephone Signals, July 2001.
- [36] IETF RFC 2401, Security Architecture for the Internet Protocol, November 1998.
- [37] IETF RFC 2402, IP Authentication Header, November 1998.
- [38] IETF RFC 2406, Type of Service in the Internet Protocol Suite, November 1998.
- [39] PacketCable 1.5 Provisioning Specification, PKT-SP-PROV1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
- [40] IETF RFC 3261, SIP: Session Initiation Protocol, June 2002.
- [41] IETF RFC 3264, An Offer/Answer Model with Session Description Protocol (SDP). June 2002.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; [http:// www.cablelabs.com](http://www.cablelabs.com).
- ITU-T Recommendations available at <http://www.itu.int>
- IETF RFCs available at <http://www.ietf.org/rfc.html>

3 TERMS AND DEFINITIONS

PacketCable specifications use the following terms:

Access Control	Limiting the flow of information from the resources of a system only to authorized persons, programs, processes, or other system resources on a network.
Active	A service flow is said to be “active” when it is permitted to forward data packets. A service flow must first be admitted before it is active.
Admitted	A service flow is said to be “admitted” when the CMTS has reserved resources (e.g., bandwidth) for it on the DOCSIS® network.
A-link	A-Links are SS7 links that interconnect STPs and either SSPs or SCPs. ‘A’ stands for “Access.”
Asymmetric Key	An encryption key or a decryption key used in public key cryptography, where encryption and decryption keys are always distinct.
Audio Server	An Audio Server plays informational announcements in PacketCable network. Media announcements are needed for communications that do not complete and to provide enhanced information services to the user. The component parts of Audio Server services are Media Players and Media Player Controllers.
Authentication	The process of verifying the claimed identity of an entity to another entity.
Authenticity	The ability to ensure that the given information is without modification or forgery and was in fact produced by the entity that claims to have given the information.
Authorization	The act of giving access to a service or device if one has permission to have the access.
Cipher	An algorithm that transforms data between plaintext and ciphertext.
Ciphersuite	A set which must contain both an encryption algorithm and a message authentication algorithm (e.g., a MAC or an HMAC). In general, it may also contain a key-management algorithm, which does not apply in the context of PacketCable.
Ciphertext	The (encrypted) message output from a cryptographic algorithm that is in a format that is unintelligible.
Cleartext	The original (unencrypted) state of a message or data. Also called plaintext.
Confidentiality	A way to ensure that information is not disclosed to anyone other than the intended parties. Information is encrypted to provide confidentiality. Also known as privacy.
Cryptanalysis	The process of recovering the plaintext of a message or the encryption key without access to the key.
Cryptographic algorithm	An algorithm used to transfer text between plaintext and ciphertext.
Decipherment	A procedure applied to ciphertext to translate it into plaintext.
Decryption	A procedure applied to ciphertext to translate it into plaintext.
Decryption key	The key in the cryptographic algorithm to translate the ciphertext to plaintext.
Digital certificate	A binding between an entity’s public key and one or more attributes relating to its identity, also known as a public key certificate.
Digital signature	A data value generated by a public-key algorithm based on the contents of a block of data and a private key, yielding an individualized cryptographic checksum.
Downstream	The direction from the headend toward the subscriber location.
Encipherment	A method used to translate plaintext into ciphertext.
Encryption	A method used to translate plaintext into ciphertext.
Encryption Key	The key used in a cryptographic algorithm to translate the plaintext to ciphertext.
Endpoint	A Terminal, Gateway or Multipoint Conference Unit (MCU).

Errored Second	Any 1-second interval containing at least one bit error.
Event Message	A message capturing a single portion of a connection.
F-link	F-Links are SS7 links that directly connect two SS7 end points, such as two SSPs. 'F' stands for "Fully Associated."
Flow [DOCSIS Flow]	(a.k.a. DOCSIS-QoS "service flow") A unidirectional sequence of packets associated with a Service ID (SID) and a QoS. Multiple multimedia streams may be carried in a single DOCSIS Flow.
Flow [IP Flow]	A unidirectional sequence of packets identified by OSI Layer 3 and Layer 4 header information. This information includes source/destination IP addresses, source/destination port numbers, protocol ID. Multiple multimedia streams may be carried in a single IP Flow.
Gateway	Devices bridging between the PacketCable IP Voice Communication world and the PSTN. Examples are the Media Gateway, which provides the bearer circuit interfaces to the PSTN and transcodes the media stream, and the Signaling Gateway, which sends and receives circuit switched network signaling to the edge of the PacketCable network.
H.323	An ITU-T recommendation for transmitting and controlling audio and video information. The H.323 recommendation requires the use of the ITU-T H.225 and ITU-T H.245 protocol for communication control between a "gateway" audio/video endpoint and a "gatekeeper" function.
Header	Protocol control information located at the beginning of a protocol data unit.
Integrity	A way to ensure that information is not modified except by those who are authorized to do so.
IntraLATA	Within a Local Access Transport Area.
Jitter	Variability in the delay of a stream of incoming packets making up a flow such as a voice communication.
Kerberos	A secret-key network authentication protocol that uses a choice of cryptographic algorithms for encryption and a centralized key database for authentication.
Key	A mathematical value input into the selected cryptographic algorithm.
Key Exchange	The swapping of public keys between entities to be used to encrypt communication between the entities.
Key Management	The process of distributing shared symmetric keys needed to run a security protocol.
Key Pair	An associated public and private key where the correspondence between the two are mathematically related, but it is computationally infeasible to derive the private key from the public key.
Keying Material	A set of cryptographic keys and their associated parameters, normally associated with a particular run of a security protocol.
Keyspace	The range of all possible values of the key for a particular cryptographic algorithm.
Latency	The time, expressed in quantity of symbols, taken for a signal element to pass through a device.
Link Encryption	Cryptography applied to data as it travels on data links between the network devices.
Network Layer	Layer 3 in the Open System Interconnection (OSI) architecture that provides network information that is independent from the lower layers.
Network Management	The functions related to the management of data across the network.
Network Management OSS	The functions related to the management of data link layer and physical layer resources and their stations across the data network supported by the hybrid fiber/coax system.

Nonce	A random value used only once that is sent in a communications protocol exchange to prevent replay attacks.
Non-Repudiation	The ability to prevent a sender from denying later that he or she sent a message or performed an action.
Off-Net Call	A communication connecting a PacketCable subscriber out to a user on the PSTN.
On-Net Call	A communication placed by one customer to another customer entirely on the PacketCable Network.
One-way Hash	A hash function that has an insignificant number of collisions upon output.
Plaintext	The original (unencrypted) state of a message or data. Also called cleartext.
Pre-shared Key	A shared secret key passed to both parties in a communication flow, using an unspecified manual or out-of-band mechanism.
Privacy	A way to ensure that information is not disclosed to any one other than the intended parties. Information is usually encrypted to provide confidentiality. Also known as confidentiality.
Private Key	The key used in public key cryptography that belongs to an individual entity and must be kept secret.
Proxy	A facility that indirectly provides some service or acts as a representative in delivering information, thereby eliminating the need for a host to support the service.
Public Key	The key used in public key cryptography that belongs to an individual entity and is distributed publicly. Other entities use this key to encrypt data to be sent to the owner of the key.
Public Key Certificate	A binding between an entity's public key and one or more attributes relating to its identity, also known as a digital certificate.
Public Key Cryptography	A procedure that uses a pair of keys, a public key and a private key, for encryption and decryption, also known as an asymmetric algorithm. A user's public key is publicly available for others to use to send a message to the owner of the key. A user's private key is kept secret and is the only key that can decrypt messages sent encrypted by the user's public key.
Root Private Key	The private signing key of the highest-level Certification Authority. It is normally used to sign public key certificates for lower-level Certification Authorities or other entities.
Root Public Key	The public key of the highest level Certification Authority, normally used to verify digital signatures generated with the corresponding root private key.
Secret Key	The cryptographic key used in a symmetric key algorithm, which results in the secrecy of the encrypted data depending solely upon keeping the key a secret, also known as a symmetric key.
Session Key	A cryptographic key intended to encrypt data for a limited period of time, typically between a pair of entities.
Signed and Sealed	An "envelope" of information which has been signed with a digital signature and sealed using encryption.
Subflow	A unidirectional flow of IP packets characterized by a single source and destination IP address and single source and destination UDP/TCP port.
Symmetric Key	The cryptographic key used in a symmetric key algorithm, which results in the secrecy of the encrypted data depending solely upon keeping the key a secret, also known as a secret key.
Systems Management	Functions in the application layer related to the management of various Open Systems Interconnection (OSI) resources and their status across all layers of the OSI architecture.

Transit Delays	The time difference between the instant at which the first bit of a Protocol Data Unit (PDU) crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.
Trunk	An analog or digital connection from a circuit switch that carries user media content and may carry voice signaling (M _F , R ₂ , etc.).
Tunnel Mode	An IPsec (ESP or AH) mode that is applied to an IP tunnel, where an outer IP packet header (of an intermediate destination) is added on top of the original, inner IP header. In this case, the ESP or AH transform treats the inner IP header as if it were part of the packet payload. When the packet reaches the intermediate destination, the tunnel terminates, and both the outer IP packet header and the IPsec ESP or AH transform are taken out.
Upstream	The direction from the subscriber location toward the headend.
X.509 certificate	A public key certificate specification developed as part of the ITU-T X.500 standards directory.

4 ABBREVIATIONS AND ACRONYMS

PacketCable specifications use the following abbreviations.

AAA	Authentication, Authorization and Accounting.
AES	Advanced Encryption Standard. A block cipher, used to encrypt the media traffic in PacketCable.
AF	Assured Forwarding. This is a DiffServ Per Hop Behavior.
AH	Authentication header. An IPsec security protocol that provides message integrity for complete IP packets, including the IP header.
AMA	Automated Message Accounting. A standard form of call detail records (CDRs) developed and administered by Bellcore (now Telcordia Technologies).
ASD	Application-Specific Data. A field in some Kerberos key management messages that carries information specific to the security protocol for which the keys are being negotiated.
AT	Access Tandem.
ATM	Asynchronous Transfer Mode. A protocol for the transmission of a variety of digital signals using uniform 53-byte cells.
BAF	Bellcore AMA Format, also known as AMA.
BCID	Billing Correlation ID.
BPI+	Baseline Privacy Plus Interface Specification. The security portion of the DOCSIS 1.1 standard that runs on the MAC layer.
CA	Certification Authority. A trusted organization that accepts certificate applications from entities, authenticates applications, issues certificates and maintains status information about certificates.
CA	Call Agent. The part of the CMS that maintains the communication state, and controls the line side of the communication.
CBC	Cipher Block Chaining mode. An option in block ciphers that combine (XOR) the previous block of ciphertext with the current block of plaintext before encrypting that block of the message.
CBR	Constant Bit Rate.
CDR	Call Detail Record. A single CDR is generated at the end of each billable activity. A single billable activity may also generate multiple CDRs.
CIC	Circuit Identification Code. In ANSI SS7, a two-octet number that uniquely identifies a DSO circuit within the scope of a single SS7 Point Code.
CID	Circuit ID (Pronounced “kid”). This uniquely identifies an ISUP DS0 circuit on a Media Gateway. It is a combination of the circuit’s SS7 gateway point code and Circuit Identification Code (CIC). The SS7 DPC is associated with the Signaling Gateway that has domain over the circuit in question.
CIF	Common Intermediate Format.
CIR	Committed Information Rate.
CM	DOCSIS Cable Modem.
CMS	Cryptographic Message Syntax.
CMS	Call Management Server. Controls the audio connections. Also called a Call Agent in MGCP/SGCP terminology. This is one example of an Application Server.
CMTS	Cable Modem Termination System. The device at a cable headend which implements the DOCSIS RFI MAC protocol and connects to CMs over an HFC network.
CMSS	Call Management Server Signaling.
Codec	COder-DECoder.

COPS	Common Open Policy Service protocol. Currently an internet draft, which describes a client/server model for supporting policy control over QoS Signaling Protocols and provisioned QoS resource management.
CoS	Class of Service. The type 4 tuple of a DOCSIS configuration file.
CRCX	Create Connection.
CSR	Customer Service Representative.
DA	Directory Assistance.
DE	Default. This is a DiffServ Per Hop Behavior.
DES	Data Encryption Standard.
DF	Delivery Function.
DHCP	Dynamic Host Configuration Protocol.
DHCP-D	DHCP Default. Network Provider DHCP Server.
DNS	Domain Name Service.
DOCSIS®	Data-Over-Cable Service Interface Specifications.
DPC	Destination Point Code. In ANSI SS7, a 3-octet number which uniquely identifies an SS7 Signaling Point, either an SSP, STP, or SCP.
DQoS	Dynamic Quality-of-Service. Assigned on the fly for each communication depending on the QoS requested.
DSA	Dynamic Service Add.
DSC	Dynamic Service Change.
DSCP	DiffServ Code Point. A field in every IP packet that identifies the DiffServ Per Hop Behavior. In IP version 4, the TOS byte is redefined to be the DSCP. In IP version 6, the Traffic Class octet is used as the DSCP.
DTMF	Dual-tone Multi Frequency (tones).
EF	Expedited Forwarding. A DiffServ Per Hop Behavior.
E-MTA	Embedded MTA. A single node that contains both an MTA and a cable modem.
EO	End Office.
ESP	IPsec Encapsulating Security Payload. Protocol that provides both IP packet encryption and optional message integrity, not covering the IP packet header.
ETSI	European Telecommunications Standards Institute.
F-link	F-Links are SS7 links that directly connect two SS7 end points, such as two SSPs. 'F' stands for "Fully Associated."
FEID	Financial Entity ID.
FGD	Feature Group D signaling.
FQDN	Fully Qualified Domain Name. Refer to IETF RFC 2821 for details.
GC	Gate Controller.
GTT	Global Title Translation.
HFC	Hybrid Fiber/Coaxial. An HFC system is a broadband bi-directional shared media transmission system using fiber trunks between the headend and the fiber nodes, and coaxial distribution from the fiber nodes to the customer locations.
HMAC	Hashed Message Authentication Code. A message authentication algorithm, based on either SHA-1 or MD5 hash and defined in IETF RFC 2104.
HTTP	Hypertext Transfer Protocol. Refer to IETF RFC 1945 and RFC 2068.
IANA	Internet Assigned Numbered Authority. See www.ietf.org for details.

IC	Inter-exchange Carrier.
IETF	Internet Engineering Task Force. A body responsible, among other things, for developing standards used on the Internet. See www.ietf.org for details.
IKE	Internet Key Exchange. A key-management mechanism used to negotiate and derive keys for SAs in IPsec.
IKE–	A notation defined to refer to the use of IKE with pre-shared keys for authentication.
IKE+	A notation defined to refer to the use of IKE with X.509 certificates for authentication.
IP	Internet Protocol. An Internet network-layer protocol.
IPsec	Internet Protocol Security. A collection of Internet standards for protecting IP packets with encryption and authentication.
ISDN	Integrated Services Digital Network.
ISTP	Internet Signaling Transport Protocol.
ISUP	ISDN User Part. A protocol within the SS7 suite of protocols that is used for call signaling within an SS7 network.
ITU	International Telecommunication Union.
ITU-T	International Telecommunication Union–Telecommunication Standardization Sector.
IVR	Interactive Voice Response system.
KDC	Key Distribution Center.
LATA	Local Access and Transport Area.
LD	Long Distance.
LIDB	Line Information Database. Contains customer information required for real-time access such as calling card personal identification numbers (PINs) for real-time validation.
LLC	Logical Link Control. The Ethernet Packet header and optional 802.1P tag which may encapsulate an IP packet. A sublayer of the Data Link Layer.
LNP	Local Number Portability. Allows a customer to retain the same number when switching from one local service provider to another.
LSSGR	LATA Switching Systems Generic Requirements.
MAC	Message Authentication Code. A fixed-length data item that is sent together with a message to ensure integrity, also known as a MIC.
MAC	Media Access Control. It is a sublayer of the Data Link Layer. It normally runs directly over the physical layer.
MC	Multipoint Controller.
MCU	Multipoint Conferencing Unit.
MD5	Message Digest 5. A one-way hash algorithm that maps variable length plaintext into fixed-length (16 byte) ciphertext.
MDCP	Media Device Control Protocol. A media gateway control specification submitted to IETF by Lucent. Now called SCTP.
MDCX	Modify Connection.
MDU	Multi-Dwelling Unit. Multiple units within the same physical building. The term is usually associated with high-rise buildings.
MEGACO	Media Gateway Control IETF working group. See www.ietf.org for details.
MF	Multi-Frequency.
MG	Media Gateway. Provides the bearer circuit interfaces to the PSTN and transcodes the media stream.

MGC	Media Gateway Controller. The overall controller function of the PSTN gateway. Receives, controls and mediates call-signaling information between the PacketCable and PSTN.
MGCP	Media Gateway Control Protocol. Protocol follow-on to SGCP. Refer to IETF 2705.
MIB	Management Information Base.
MIC	Message Integrity Code. A fixed-length data item that is sent together with a message to ensure integrity, also known as a Message Authentication Code (MAC).
MMC	Multi-Point Mixing Controller. A conferencing device for mixing media streams of multiple connections.
MSB	Most Significant Bit.
MSO	Multi-System Operator. A cable company that operates many headend locations in several cities.
MSU	Message Signal Unit.
MTA	Multimedia Terminal Adapter. Contains the interface to a physical voice device, a network interface, CODECs, and all signaling and encapsulation functions required for VoIP transport, class features signaling, and QoS signaling.
MTP	The Message Transfer Part. A set of two protocols (MTP 2, MTP 3) within the SS7 suite of protocols that are used to implement physical, data link, and network-level transport facilities within an SS7 network.
MWD	Maximum Waiting Delay.
NANP	North American Numbering Plan.
NANPNAT	North American Numbering Plan Network Address Translation.
NAT Network Layer	Network Address Translation. Layer 3 in the Open System Interconnection (OSI) architecture. This layer provides services to establish a path between open systems.
NCS	Network Call Signaling.
NPA-NXX	Numbering Plan Area (more commonly known as area code) NXX (sometimes called exchange) represents the next three numbers of a traditional phone number. The N can be any number from 2-9 and the Xs can be any number. The combination of a phone number's NPA-NXX will usually indicate the physical location of the call device. The exceptions include toll-free numbers and ported number (see LNP).
NTP	Network Time Protocol. An internet standard used for synchronizing clocks of elements distributed on an IP network.
NTSC	National Television Standards Committee. Defines the analog color television broadcast standard used today in North America.
OID	Object Identification.
OSP	Operator Service Provider.
OSS	Operations Systems Support. The back-office software used for configuration, performance, fault, accounting, and security management.
OSS-D	OSS Default. Network Provider Provisioning Server.
PAL	Phase Alternate Line. The European color television format that evolved from the American NTSC standard.
PCES	PacketCable Electronic Surveillance.
PCM	Pulse Code Modulation. A commonly employed algorithm to digitize an analog signal (such as a human voice) into a digital bit stream using simple analog-to-digital conversion techniques.
PDU	Protocol Data Unit.
PHS	Payload Header Suppression. A DOCSIS technique for compressing the Ethernet, IP, and UDP headers of RTP packets.

PKCROSS	Public-Key Cryptography for Cross-Realm Authentication. Utilizes PKINIT for establishing the inter-realm keys and associated inter-realm policies to be applied in issuing cross-realm service tickets between realms and domains in support of Intradomain and Interdomain CMS-to-CMS signaling (CMSS).
PKCS	Public-Key Cryptography Standards. Published by RSA Data Security Inc. These Standards describe how to use public key cryptography in a reliable, secure and interoperable way.
PKI	Public-Key Infrastructure. A process for issuing public key certificates, which includes standards, Certification Authorities, communication between authorities and protocols for managing certification processes.
PKINIT	Public-Key Cryptography for Initial Authentication. The extension to the Kerberos protocol that provides a method for using public-key cryptography during initial authentication.
PSC	Payload Service Class Table, a MIB table that maps RTP payload Type to a Service Class Name.
PSFR	Provisioned Service Flow Reference. An SFR that appears in the DOCSIS configuration file.
PSTN	Public Switched Telephone Network.
QCIF	Quarter Common Intermediate Format.
QoS	Quality of Service. Guarantees network bandwidth and availability for applications.
RADIUS	Remote Authentication Dial-In User Service. An internet protocol (IETF RFC 2865 and RFC 2866) originally designed for allowing users dial-in access to the internet through remote servers. Its flexible design has allowed it to be extended well beyond its original intended use.
RAS	Registration, Admissions and Status. RAS Channel is an unreliable channel used to convey the RAS messages and bandwidth changes between two H.323 entities.
RC4	Rivest Cipher 4. A variable length stream cipher. Optionally used to encrypt the media traffic in PacketCable.
RFC	Request for Comments. Technical policy documents approved by the IETF which are available on the World Wide Web at http://www.ietf.cnri.reston.va.us/rfc.html .
RFI	The DOCSIS Radio Frequency Interface specification.
RJ-11	Registered Jack-11. A standard 4-pin modular connector commonly used in the United States for connecting a phone unit into a wall jack.
RKS	Record Keeping Server. The device, which collects and correlates the various Event Messages.
RSA	A public-key, or asymmetric, cryptographic algorithm used to provide authentication and encryption services. RSA stands for the three inventors of the algorithm; Rivest, Shamir, Adleman.
RSA Key Pair	A public/private key pair created for use with the RSA cryptographic algorithm.
RSVP	Resource Reservation Protocol.
RTCP	Real-Time Control Protocol.
RTO	Retransmission Timeout.
RTP	Real-time Transport Protocol. A protocol for encapsulating encoded voice and video streams. Refer to IETF RFC 1889.
SA	Security Association. A one-way relationship between sender and receiver offering security services on the communication flow.
SAID	Security Association Identifier. Uniquely identifies SAs in the DOCSIS Baseline Privacy Plus Interface (BPI+) security protocol.
SCCP	Signaling Connection Control Part. A protocol within the SS7 suite of protocols that provides two functions in addition to those provided within MTP. The first function is the ability to address applications within a signaling point. The second function is Global Title Translation.

SCP	Service Control Point. A Signaling Point within the SS7 network, identifiable by a Destination Point Code that provides database services to the network.
SCTP	Stream Control Transmission Protocol.
SDP	Session Description Protocol.
SDU	Service Data Unit. Information delivered as a unit between peer service access points.
SF	Service Flow. A unidirectional flow of packets on the RF interface of a DOCSIS system.
SFID	Service Flow ID. A 32-bit integer assigned by the CMTS to each DOCSIS Service Flow defined within a DOCSIS RF MAC domain. SFIDs are considered to be in either the upstream direction (USFID) or downstream direction (DSFID). Upstream Service Flow IDs and Downstream Service Flow IDs are allocated from the same SFID number space.
SFR	Service Flow Reference. A 16-bit message element used within the DOCSIS TLV parameters of Configuration Files and Dynamic Service messages to temporarily identify a defined Service Flow. The CMTS assigns a permanent SFID to each SFR of a message.
SG	Signaling Gateway. An SG is a signaling agent that receives/sends SCN native signaling at the edge of the IP network. In particular, the SS7 SG function translates variants ISUP and TCAP in an SS7-Internet Gateway to a common version of ISUP and TCAP.
SGCP	Simple Gateway Control Protocol. Earlier draft of MGCP.
SHA – 1	Secure Hash Algorithm 1. A one-way hash algorithm.
SID	Service ID. A 14-bit number assigned by a CMTS to identify an upstream virtual circuit. Each SID separately requests and is granted the right to use upstream bandwidth.
SIP	Session Initiation Protocol. An application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants.
SIP+	Session Initiation Protocol Plus. An extension to SIP.
S-MTA	Standalone MTA. A single node that contains an MTA and a non-DOCSIS MAC (e.g., ethernet).
SNMP	Simple Network Management Protocol.
SOHO	Small Office/Home Office.
SS7	Signaling System number 7. An architecture and set of protocols for performing out-of-band call signaling with a telephone network.
SSP	Service Switching Point. SSPs are points within the SS7 network that terminate SS7 signaling links and also originate, terminate, or tandem switch calls.
STP	Signal Transfer Point. A node within an SS7 network that routes signaling messages based on their destination address. This is essentially a packet switch for SS7. It may also perform additional routing services such as Global Title Translation.
TCAP	Transaction Capabilities Application Protocol. A protocol within the SS7 stack that is used for performing remote database transactions with a Signaling Control Point.
TCP	Transmission Control Protocol.
TD	Timeout for Disconnect.
TFTP	Trivial File Transfer Protocol.
TFTP-D	Default – Trivial File Transfer Protocol.
TGS	Ticket Granting Server. A sub-system of the KDC used to grant Kerberos tickets.
TGW	Telephony Gateway.
TIPHON	Telecommunications and Internet Protocol Harmonization Over Network.
TLV	Type-Length-Value. A tuple within a DOCSIS configuration file.
TN	Telephone Number.

ToD	Time-of-Day Server.
TOS	Type of Service. An 8-bit field of every IP version 4 packet. In a DiffServ domain, the TOS byte is treated as the DiffServ Code Point, or DSCP.
TSG	Trunk Subgroup.
UDP	User Datagram Protocol. A connectionless protocol built upon Internet Protocol (IP).
VAD	Voice Activity Detection.
UDPTL	UDP Transport Layer for fax transmission based on ITU T.38 protocol
VBR	Variable Bit Rate.
VoIP	Voice-over-IP.

5 SCOPE

This specification describes a profile of the Media Gateway Control Protocol (MGCP) for PacketCable embedded clients, which we will refer to as the PacketCable Network-based Call Signaling (NCS) protocol. MGCP is a call signaling protocol for use in a centralized call control architecture, and assumes relatively simple client devices. The call signaling protocol is one layer of the overall PacketCable suite of specifications and relies upon companion protocol specifications to provide complete end-to-end PacketCable functionality. The scope of NCS is currently only embedded Voice-Over-IP client devices in a PacketCable environment and the NCS profile has therefore simplified and in some cases modified the base MGCP 1.0 protocol accordingly. Support for video will be added in a later version of this document.

This document describes a PacketCable profile of an application programming interface (MGCI), and a corresponding 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 gateways and is handled by external call control elements.

This document is based on the Media Gateway Control Protocol (MGCP) 1.0 RFC [1], which was the result of a merge of the IETF draft of Simple Gateway Control Protocol, and the IETF draft of IP Device Control (IPDC) family of protocols, as well as input generated by the PacketCable embedded client signaling team that developed this profile.

This document, which defines the PacketCable NCS Protocol specification, constitutes a document that is independent of MGCP in order to provide a stable reference document, while meeting current time-to-market demands for such a reference. It is the intent of this document to be as closely aligned with MGCP as possible for the PacketCable environment, in order to avoid developing multiple protocols to solve the same problem. This goal has been, and continues to be, pursued through cooperation with the authors of the MGCP specification. The NCS profile of MGCP, however, is strictly and solely defined by the contents of this document.

This NCS profile of MGCP, referred to as the Network-based Call Signaling Protocol 1.0, NCS 1.0, the NCS profile, or simply NCS in this document, has been modified from the MGCP 1.0 RFC in the following ways:

- *The NCS protocol only aims at supporting PacketCable-embedded clients:* The NCS protocol supports embedded clients as defined by PacketCable. Functionality present in the MGCP 1.0 protocol, which was superfluous to NCS, has been removed.
- *The NCS protocol contains extensions and modifications to MGCP:* PacketCable-specific requirements have been addressed in NCS, which has resulted in a couple of minor extensions and modifications to MGCP. However, the MGCP architecture, and all of the MGCP constructs relevant to embedded clients, are preserved in NCS.
- *The NCS protocol contains minor simplifications from MGCP 1.0:* Where several choices were available, and not necessarily needed for an embedded client in the PacketCable environment, some simplifications have been made for embedded-client implementations.

Although MGCP is not NCS, and NCS is not MGCP, the names MGCP and NCS will be used interchangeably in this document since this document is based on MGCP. Unless otherwise stated or inferred by context, the word MGCP shall be taken to mean the NCS profile of MGCP in this document.

The document is structured in the following main sections:

- The introduction presents the basic assumptions and the relation to other protocols such as H.323, RTSP, SAP, and SIP.
- The interface section presents a conceptual overview of the NCS profile, presenting the naming conventions, the usage of the session description protocol (SDP), and the procedures that compose NCS:
 - Notification Request,
 - Notify,
 - Create Connection,
 - Modify Connection,
 - Delete Connection,
 - Audit Endpoint,
 - Audit Connection, and
 - Restart In Progress.
- Each of the commands is presented as an example API with the name of the command, the parameters it can take and return, as well as the semantics of each of these. The actual encoding of the command and its parameters is shown in the protocol description section. The section concludes with a description of how reliability and race conditions are handled.
- The protocol description section presents the actual NCS encodings of the commands and parameters, which are based on simple text formats. The transmission procedure over UDP is specified as well.
- The security section presents the security of NCS.
- The appendices contain definitions of event packages, mode interactions, example command encodings, and a complete example call flow.

6 INTRODUCTION

This 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 analog (RJ11) access lines to a voice-over-IP (VoIP) network.
- Optionally, one or more video lines to a VoIP network.

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., analog 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 this document does not define a mechanism for synchronizing Call Agents, although future PacketCable 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 analog 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.

6.1 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 below:




Figure 1 - Relation to H.323 Standards

In the MGCP model, the gateways focus on the audio signal translation function, while the Call Agent handles the signaling and call processing functions. As a consequence, the Call Agent implements the "signaling" 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.

6.2 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 multi-media applications. These other specifications include:

- the session description protocol (SDP) [4],
- the session announcement protocol (SAP) [5],
- the session initiation protocol (SIP) [40],
- the real-time streaming protocol (RTSP) [7].

The latter three specifications are, in fact, alternative signaling 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. The Call Agent provides for signaling conversion, as illustrated in Figure 2 below:




Figure 2 - Relation to IETF Standards

The SDP standard has a pivotal status in this architecture. In the following description it is used to carry session descriptions in MGCP.

6.3 Relation to RFC 3435 and ABNF Grammar

RFC 3435 includes a formal description of the MGCP protocol syntax following the "Augmented BNF for Syntax Specifications". This formal description is referenced by developers for the creation of interoperable devices. A copy of the MGCP protocol syntax, annotated and edited to indicate its applicability to PacketCable specifications, is provided in Appendix G. Adherence to these guidelines can improve interoperability by minimizing failures caused by different interpretations of syntax and grammar.

7 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.

7.1 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.

7.1.1 Endpoint Names

Endpoint names, also known as (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 [16] 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 [14] 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 analog 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". Unless specified otherwise, this refers to all endpoints configured for service, regardless of their actual service state, i.e., in-service or out-of-service.
- 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". Unless specified otherwise, this only refers to endpoints that are in-service.
- 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. Note that since the domain name portion is part of the endpoint identifier, different forms or different values referring to the same entity are not freely interchangeable. Following a restart, the most recently supplied form and value MUST always be used.

7.1.1.1 Embedded Client Endpoint Names

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

Embedded clients will support the following two endpoint-types:

- Analog telephone: The Analog Telephone is represented as an analog access line (aaln). This is basically the equivalent of an analog telephone line as known in the PSTN (see [10]).
- Video: The details of the video device-type will be provided in a future version of this document.

7.1.1.1.1 Analog Access Line Endpoints

In addition to the naming conventions specified above, local endpoint names for endpoints of type "analog 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 can 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 analog access lines supported by the embedded client in question. The number thus identifies a specific analog 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 same wildcard character is then assumed for term2, i.e., "*" and "\$" is equivalent to respectively "*/*" and "\$/\$".

Example analog access line local endpoint names could thus be:

aaln/1	The first analog access line on the embedded client in question.
aaln/2	The second analog access line on the embedded client in question.
aaln/\$	Any analog access line on the embedded client in question.
aaln/*	All analog 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.

7.1.1.1.2 Video Endpoints

Details on video endpoints will be provided in a future version of this document.

7.1.2 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 characters 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 this 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 may all 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.

7.1.3 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 minutes, 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.

7.1.4 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 failover 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 explicitly¹, 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."

Section 7.4 contains a more detailed description of reliability and failover.

7.1.5 Digit Maps

The Call Agent can ask the gateway to collect digits dialed by the user. This facility is intended to be used for analog 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 section.

An alternative procedure involves the gateway notifying the Call Agent of the dialed digits as soon as they are dialed, a.k.a., overlap sending. However, such a procedure generates a large number of interactions. It is preferable to accumulate the dialed 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 numbers illustrated in Figure 3 below:

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

Figure 3 - Dialed Digits - Examples

The solution to this problem is to load the gateway with a digit map that corresponds to the dial plan. 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|8xxxxxxx|#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
```

¹ This could happen as a result of specifying an empty NotifiedEntity parameter.

	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 string². 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 doesn't 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 Agent³ and clear the "current dial string".. A match, in this specification, can be either a "perfect match," exactly matching one of the specified alternatives, 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 map⁴, 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.
- 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 seconds and the default value for T_{crit} is 4 seconds. The provisioning process may alter both of these.

Appendix A contains additional detail and an example on use of timer T.

The end-points MUST support at least 2048 Bytes of digit map 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. Digit Maps, when provided by the Call Agent, MUST be as defined in this section.

² For instance, "123T" and "123[1-2T5]" satisfy that rule, but "12T3" does not.

³ The list of digits may include other events as well. See Section 7.4.3.1

⁴ Technically speaking with the "accumulate according to digit map" action.

7.1.6 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 analog 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 analog access line with the line package (package name "L") being the default package, the following two event names are considered equal:

L/dl dial-tone in the line package for an analog access line.

dl dial-tone in the line package (default) for an analog access line.

This document defines the packages for embedded client types listed in Table 1.

The CMS **MUST** support the "L", "FXR" and "XRM" packages. The CMS **MAY** support the "VBD", "GPMD" and "FM" packages defined in Appendix A.5. The MTA **MUST** support the "L", "FXR" and "XRM" packages. The MTA **MAY** support the "VBD", "GPMD" and "FM" packages defined in Appendix A.5.

Table 1 - Packages for Embedded Client End-point Types

Endpoint-Type	Package	Package Name	Default Package	1st Release that Introduced Package Support Requirement
Analog Access Line	Line	L	Yes	1.0
Analog Access Line	FAX	FXR	No	1.5
Analog Access Line	VoIP Metrics	XRM	No	1.5
Analog Access Line	VoiceBandData	VBD	No	1.5
Analog Access Line	General Purpose Media Descriptor	GPMD	No	1.5
Analog Access Line	Media Format Parameter Package	FM	No	1.5

Additional package names and event codes may be defined by and/or registered with PacketCable. Any change to the packages defined in this 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 Section 7.3.1) and if they contain auditable event-states (e.g., off-hook, see Section 7.3.8.1). 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 Section 7.3.1.

In addition to PacketCable packages, implementers **MAY** gain experience by defining experimental packages. The package name of experimental packages **MUST** begin with the two characters "x-" or "X-"; PacketCable **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:

L/all refers to all events in the line package for an analog access line.

*/all for an analog 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:

L/rt@0A3F58

Should the connection be deleted while an event or signal is being detected or applied on it, that particular event detection or signal generation simply MUST stop. Depending on the signal, the endpoint SHOULD generate a failure, i.e., if the signal type is TO, the operation failure event will be generated, since the connection associated with the signal was deleted prior to the signal timing out. Notification action associated with reporting the failure must conform to the notify operations as defined for Notification Request handling (section 7.3.1).

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:

L/ma@*

However, when the event is actually observed, the gateway MUST include the name of the specific connection on which the event occurred. 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:

L/rt@\$

When processing a command using the "current connection" wildcard, the "\$" wildcard character MUST be expanded by the gateway to the value of the current connection. If a subsequent command either explicitly (e.g., by auditing) or implicitly (e.g., by persistence) refers to such an event, the expanded value MUST be used by the gateway. In other words, the "current connection" wildcard is expanded once, which is at the initial processing of the command in which it was explicitly included. 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). However, the use of the "all packages" and "all events" wildcards is strongly discouraged.

Call Agents must be able to operate in an environment where some endpoints do not support all packages. An endpoint that receives a command referencing a package it does not support will respond with error 518 (Unsupported or Unknown Package). On receiving this error response, the Call Agent could try the command again without the package parameter, although in the case where the original command contained parameters for multiple packages, the Call Agent would have no way of knowing which specific package(s) to exclude. The Call Agent can also use the AuditEndpoint command to determine the set of packages supported by an endpoint.

7.2 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 this specification, SDP descriptions **MUST** follow the conventions delineated in the session description protocol (SDP), which is now an IETF-Proposed Standard documented in [4]. In addition, all call agents and gateways **MUST** ignore any SDP parameters, attributes, or fields that are not understood by the call agent or gateway.

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.

SDP allows for description of real time fax using media type "image". The NCS profile will support setting of fax connections using media type "image".

7.3 Gateway Control Functions

This section 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 section 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 behavior 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 (Section 8). Embedded clients and Call Agents **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 is 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 this specification.
- 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 analog 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. Flow of media on each connection is controlled 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" (netwtest). The "mode" parameter determines if media packets can be sent and/or received on the connection. RTCP is independent of the connection mode; for more details see the Codec specification [11].

Handling of media received from the endpoint is determined by the mode parameter:

- Media originating from the endpoint will be sent on all the connections for that endpoint whose mode is either "send only", "send/receive", "conference", or "replicate".

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

- Media received in data packets through connections in "inactive", "sendonly", or "replicate" mode is discarded.
- Media received in data packets through connections in "receive only", "conference", or "send/receive" mode is mixed together and then sent to the endpoint.
- In addition to being sent to the endpoint, media received in data packets through connections in "conference" mode is 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., are outside the scope of this document.
- In addition to the media received from the endpoint, media sent to the endpoint is mixed and transmitted over all the other connections for the endpoint whose mode is "replicate". This SHOULD include media generated by signals applied to the endpoint.
- Media 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 analog access lines, the signal will be an audio signal, and the signal MUST NOT be passed on to a telephone connected to the analog 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 G.711 encoding⁵. Support of the "conference" mode is optional; all other connection modes MUST be supported. Please refer to Appendix B for illustrations of mode interactions.

⁵ The "replicate" connection can, e.g., be used to support "busy line verification" with minimal resource impact on the embedded client.

7.3.1 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 accordingly⁶.

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 Section 7.1.1. The "any of" wildcard MUST NOT be used. An embedded client that receives a NotificationRequest with the "any of" wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response. The "all of" wildcard MUST be supported for NotificationRequests with each of RequestedEvents, SignalsRequest, DigitMap, and DetectEvents being either empty or omitted. For simplicity, some gateways may choose to not support the "all-of" wildcard for NotificationRequests where one or more of these parameters is neither empty nor omitted. Such gateways shall respond with error code 503 if they receive an "all-of" wildcarded NotificationRequest which they are unable to process for this reason.

NotifiedEntity is an optional parameter that specifies a new "notified entity" for the endpoint. When used, the entire Call Agent name MUST be specified which includes both the local name and domain name - even if a bracketed IP address⁷ is used for the domain name. See Sections 7.1.1 and 7.1.4 for more information. If, however, only the domain name is provided, the MTA SHOULD use the domain name as the Call Agent ID.

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 signals⁸:

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

⁶ The new instruction would be a ModifyConnection command.

⁷ The use of an IP address in the NotifiedEntity is not permitted by the security specification [13]. When implementing the security specification [13], an absolute domain name (including the hostname) must be used.

⁸ Please refer to Appendix A for a complete list of signals.

Signals are divided into different types depending upon their behavior:

- **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, an OO signal MUST NOT be turned off until explicitly instructed to by the Call Agent; OO signals will be off following restart of the endpoint. A missing or empty SignalRequest parameter or a SignalRequest parameter that omits a particular signal is not an explicit instruction for an OO signal; it will not change the MTA state for an OO signal.
- **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 Appendix A for further definition of this event). A TO signal could be "ringback" timing out after 180 seconds. If an event occurs prior to the 180 seconds, the signal will, by default, be stopped⁹. 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 out¹⁰. 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 time out¹⁰.
- **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 behavior.

When a signal is applied to a connection that has received a RemoteConnectionDescriptor (see Section 7.3.3), the media stream generated by that signal MUST be forwarded on the connection for all modes except "network loopback". 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 omission of the SignalRequests parameter is interpreted as an empty SignalRequests list.

The currently defined signals can be found in Appendix A.

⁹ The "Keep signal(s) active" action may override this behavior.

¹⁰ If parameters were passed to the signal, the parameters will not be reported.

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:

- fax tones,
- modem tones,
- 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),
- flash hook (occurring in classic telephone sets when the user briefly presses the hook that holds the handset),
- DTMF digits (or pulse digits).

The currently defined events can be found in Appendix A.

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 action¹¹. 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 performed¹². Persistent events are identified as such through their definition – see Appendix A.

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. The omission of the RequestedEvents parameter is interpreted as an empty RequestedEvents list.

¹¹ Thus, the RequestIdentifier will be the RequestIdentifier of the current NotificationRequest.

¹² 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.

More than one action can be specified for an event, although a given action can not appear more than once for a given event. Table 2 below specifies the legal combinations of actions:

Table 2 - Legal Combinations of Events and Actions

	Notify	Accumulate	Accumulate according to digit map	Ignore	Keep Signal(s) Active	Embedded Notification Request	Embedded Modify Connection
Notify	-	-	-	-	✓	✓	✓
Accumulate	-	-	-	-	✓	✓	✓
Accumulate according to digit map	-	-	-	-	✓	-	✓
Ignore	-	-	-	-	✓	-	✓
Keep Signal(s) active	✓	✓	✓	✓	-	✓	✓
Embedded Notification Request	✓	✓	-	-	✓	-	✓
Embedded Modify Connection	✓	✓	✓	✓	✓	✓	-

Note: The "Embedded Notification Request" action can only be combined with the "Notify" action, if the gateway is allowed to issue more than one Notify command per Notification request.

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.

When a stimulus that triggers multiple requested events is detected (e.g., fax tone is the stimulus for both FXR/gwfax(start), and L/ft), the gateway MUST generate only one of the events (namely the most preferred event out of the multiple requested events that triggered) based on the following precedence rules:

1. Events included in a RequestedEvents list are prioritized from left to right, with the most preferred event listed to the left.
2. Persistent events that are not included in a RequestedEvents lists are less preferred than events (persistent or not) that are included in a RequestedEvents list. There is no well-defined preference order among persistent events that are not included in a RequestedEvents list.

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 Section 7.1.5. 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 Section 7.4.3.1 for further details.

The signals being applied by the SignalRequests is 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 mandates a "ringing" signal and the event request asks to look for an "off-hook" event, the ringing shall, by default, stop as soon as the gateway detects an off-hook event. If the event request did not ask to look for an "off-hook" event, the ringing would stop anyway since off-hook is a persistent event and therefore 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 behavior would be to stop all active time-out signals when the first digit¹³ is accumulated—it is irrelevant to this synchronization if the accumulated digit results in a match, mismatch, or partial matching to the digit map.

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 Section 7.4.3.1). 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 command¹⁴. 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 Appendix 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 Appendix 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.

¹³ Digit as defined in digit maps, i.e., including asterisk, timer, etc.

¹⁴ 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.

Section 7.4.3.1 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 Appendix 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 a few exceptions to this rule, e.g., fax and modem tones, which can be detected but can not be signaled. 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 Section 7.4.3.1), 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) and "loop" mode is mandatory. 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 Section 7.4.3.1.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see Section 7.5) optionally followed by commentary.

7.3.2 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 Section 7.1.1. 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. A Call Agent that receives a Notification with wildcard convention **MUST** return an error (the error returned **SHOULD** be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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. Note that the MTA **MAY** include only the absolute domain name (including the hostname) of its NotifiedEntity if only the absolute domain name was received in the triggering NotificationRequest. The CMS **SHOULD** accept the value in this case. 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 Section 7.5) optionally followed by commentary.

7.3.3 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 under-specified 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 **MUST** be returned in the Specific EndpointId parameter of the response only if the command is successful. In this case, the endpoint assigned **MUST** be in service and **MUST NOT** already have any connections on it. The "all" wildcard convention **MUST NOT** be used. An embedded client that receives a CreateConnection with the "all" wildcard convention **MUST** return an error (the error returned **SHOULD** be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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(s) **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. See Section 7.7 for details on the codec selection process. Note "encoding method" includes audio, image and video encodings.
 - The endpoint **MUST** additionally indicate which of the remaining compression algorithms it is willing to support as alternatives - see Section 8.4.1 for details.
 - A list of permissible encoding methods is specified in the PacketCable Audio/Video Codecs Specification [11]. The literal names defined in Table 4 of the Audio/Video Codes Specification [11] **MUST** be used. Unknown compression algorithms **SHOULD** be ignored if they are received.

- **Packetization Period:** A single packetization period in milliseconds, MAY be specified with exactly one decimal value. If this specifier is used, then the same packetization period MUST be used for all encoding methods allowed by the LocalConnectionOptions. Note that if no encoding method field is specified in the LCO, the MTA MUST NOT choose an encoding method with a packetization period that differs from that specified here. If different packetization periods for different encoding methods is desired, then this field MUST NOT be used. The value pertains to both media sent and received. Note that only the valid packetization period in conjunction with the associated encoding method or methods are to be used by the MTA. A list of permissible packetization periods is specified in the PacketCable Audio/Video Codecs Specification [11]. This specifier MUST NOT be supplied in the same LCO as the Multiple Packetization Period field. An MTA MUST return an error (error code 524 – inconsistency in LocalConnectionOptions) when it receives an LCO with both the Packetization Period and Multiple Packetization Period fields.
- **Multiple Packetization Period:** A list of packetization periods in milliseconds, MAY be specified if, and only if, the Encoding Method field is included. When specified, the multiple packetization period in milliseconds MUST contain exactly one decimal value or a hyphen for each entry in the encoding method field included in the LocalConnectionOptions. This applies even if several of the encoding methods have the same value. The first entry in the list MUST be a decimal number. When a hyphen is used, the encoding method in question MUST use the same packetization period as one of the other entries in the list that actually contains a decimal number, and furthermore the encoding method in question MUST NOT consume any more bandwidth than that other entry. This can for example be used for non-voice codecs (e.g., telephone-event or comfort noise) that use the same packetization period as the voice codec with which they are being used. Successive entries in the list of packetization periods MUST be ordered identically to the corresponding encoding methods. The values pertain to both media sent and received. Note that the MTA MUST NOT choose a codec with a packetization period that differs from that specified here. Note that only the valid packetization period in conjunction with the associated encoding method is to be used by the MTA. A list of permissible packetization periods is specified in the PacketCable Audio/Video Codecs Specification [11]. This specifier MUST NOT be supplied in the same LCO as the Packetization Period field. An MTA MUST return an error (error code 524 – inconsistency in LocalConnectionOptions) under the following conditions:
 - When it receives an LCO with both the Packetization Period and Multiple Packetization Period fields.
 - When it receives an LCO where the number of codecs specified in the Encoding Method field is different from the number of elements in the Multiple Packetization Period field.
- **Echo Cancellation:** Whether echo cancellation should be used initially on the line side or not¹⁵. 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 initially. The embedded client SHOULD subsequently enable or disable echo cancellation in accordance with ITU-T recommendation V.8 and V.25 when voiceband data is detected. For re-enabling echo cancellation, see e.g., G.168. Following termination of voiceband data, the handling of echo cancellation MUST revert to the current value of the echo cancellation parameter. It is RECOMMENDED that echo cancellation handling is left to the embedded client rather than having this parameter specified by the Call Agent.
- **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.¹⁶ The parameter is optional. When the parameter is omitted, a default value of 0x00 (unless provisioned otherwise) MUST be used. When this parameter is present and valid, the value provided MUST be used by the embedded client to populate the differentiated services codepoint (DSCP) parameter in the IP header (see [26] for more information on the DSCP). The parameter value MUST be 0x00 or MUST be a multiple of four within the range 0x01 to 0xFF (bits 6 and 7, the ECN bits, are reserved and hence must be set to "00"). An MTA MUST return an error (error code 532 – Unsupported value(s) in LocalConnectionOptions)

¹⁵ Echo cancellation on the packet side is not supported.

¹⁶ For RTP media streams, this type of service value applies only to the RTP media packets, not the RTCP packets.

when it receives an invalid value. The left-most "bit" in the parameter corresponds to the most significant bit in the IP header.

- **Silence Suppression:** The embedded client may perform voice activity detection, and avoid sending packets during periods of silence. However, it is necessary, for certain calls (e.g., modem calls) to disable silence suppression. 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 "off". If the value is "on", upon detecting voiceband data, the endpoint **SHOULD** disable silence suppression. Following termination of voiceband data, the handling of silence suppression **MUST** revert to the current value of the silence suppression parameter.

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

- **D-QoS GateID:** The GateID for the gate that has been setup at the CMTS. The GateID 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 **MUST** be performed for this command. The absence of this parameter indicates that D-QoS **MUST NOT** 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:

"Send Reserve"	Resources are reserved in the send direction only.
"Receive Reserve"	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.
"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 omitted, 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. Please see the PacketCable Dynamic Quality of Service Specification [12] 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 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 omitted, resources **MUST** by default be committed based on the connection mode as specified in Table 3 below:

Table 3 - Default Resource Reservation Values

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 committed¹⁷, 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.

- **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. However, this parameter **MUST** be used for resource reservation by the embedded client if provided by the Call Agent.
- **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 known¹⁸. When a RemoteConnectionDescriptor has been provided, the parameter is ignored.

The following LocalConnectionOptions fields are used to support the PacketCable 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. The endpoint **MUST** choose exactly one of the ciphersuites according to the rules described in the PacketCable Security Specification [13]. The endpoint **SHOULD** additionally indicate which of the remaining ciphersuites it is willing to support as alternatives (see Section 8.4.1 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 PacketCable Security Specification [13]. The RTP ciphersuite parameter applies to RTP media streams only. If the CMS includes LocalConnectionOptions that require use of non-RTP media only (as for example T.38 fax relay using UDPTL), the RTP ciphersuite parameter **MUST NOT** be included. If an LCO (local connection options) allows both RTP and non-RTP media and an RTP ciphersuite parameter is included, it applies to the RTP media only. In all cases, if the resulting media stream on the connection is not an RTP media stream (as for example T.38 fax relay using UDPTL), the RTP ciphersuite parameter **MUST** be ignored, i.e., RTP security is not being used, and RTP security parameters are not included in LocalConnectionDescriptor.
- **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. The endpoint **MUST** choose exactly one of the ciphersuites according to the rules described in the PacketCable Security Specification [13]. The endpoint **SHOULD** additionally indicate which of the remaining ciphersuites it is willing to support as alternatives. (See Section 8.4.1 for details). Each ciphersuite is represented as an ASCII string consisting of two substrings separated by a slash ("/"), where the first substring identifies the authentication algorithm, and

¹⁷ This is not possible for the CreateConnection command but is noted here for completeness. It is possible for the ModifyConnection command however (see Section 7.3.4).

¹⁸ Note that this will enable certain theft-of-service scenarios. See PacketCable Dynamic Quality of Service Specification [12] for details.

the second substring identifies the encryption algorithm. A list of permissible ciphersuites is specified in the PacketCable Security Specification [13]. The RTCP ciphersuite parameter applies to RTCP for RTP media streams only. If the CMS includes LocalConnectionOptions that require use of non-RTP media only (as for example T.38 fax relay using UDPTL), the RTCP ciphersuite parameter MUST NOT be included. If an LCO allows both RTP and non-RTP media and an RTCP ciphersuite parameter is included, it applies to RTCP for the RTP media only. In all cases, if the resulting media stream on the connection is not an RTP media stream (as for example T.38 fax relay using UDPTL), the RTCP ciphersuite parameter MUST be ignored, i.e., RTCP security is not being used, and RTCP security parameters are not included in LocalConnectionDescriptor.

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. Section 8.4 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 Section 7.3. 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", or "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. The ConnectionId MUST be included with any provisional or successful response to a CreateConnection command. The ConnectionId MUST NOT be included when any error response is returned, and the connection was not created.

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. Section 8.4 details the supported use of SDP in the NCS profile. The LocalConnectionDescriptor MUST be included with any provisional or successful response to a CreateConnection command. The LocalConnectionDescriptor MUST NOT be included when any error response is returned, and the connection was not created.

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 behavior 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 RemoteConnectionDescriptor for the connection, the connection can by definition not be in any of the modes "send only", "send/receive", "replicate", or "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 canceling signals and of stop looking for events. Note that if the RequestedEvents and SignalRequests parameters are omitted, then the corresponding lists are considered empty only if a RequestIdentifier parameter is included.

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 Section 7.5) optionally followed by commentary.

ResourceID is a D-QoS parameter that MAY be returned by the gateway. Note that this parameter MUST be returned by the MTA when D-QoS is to be performed as instructed by the Call Agent. When a successful D-QoS resource reservation is made, the ResourceID provides a handle for the resources reserved. The ResourceID MUST NOT be included when any error response is returned, and the connection was not created.

7.3.4 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. An embedded client that receives a ModifyConnection with a wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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 MUST then be used by the MTA.
- **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 Appendix C). 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 MUST be used by the embedded client for resource reservation and replaces the ResourceID kept 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. The LocalConnectionDescriptor MUST NOT be included when any error response is returned, and the connection was not modified. 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 change¹⁹.

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 Agent²⁰. 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 IP address specified as well. Any media streams received from any other addresses SHOULD be discarded. The PacketCable Security Specification [13] should be consulted for additional security requirements.

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. Note that absence of the RequestedEvents and SignalRequests parameters is interpreted as an empty list only if a RequestIdentifier parameter is included.

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.

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

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see Section 7.5) optionally followed by commentary.

ResourceID is a D-QoS parameter that MUST be 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. The Resource ID MUST NOT be included when any error response is returned, and the connection was not modified.

¹⁹ 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.

²⁰ For instance, if media needs to traverse a firewall.

7.3.5 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. An embedded client that receives a Delete Connection with a wildcard convention **MUST** return an error (the error returned **SHOULD** be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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. The connection parameters **MUST** only be returned if the command is successful and the connection is deleted. 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 **MUST** be based on the same information provided via the RTCP mechanism.
- **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 **MUST** be based on the same information provided via the RTCP mechanism.
- **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. All received packets **MUST** be counted independent of the connection mode or any type of processing error, e.g., authentication failure.
- **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. All received octets **MUST** be counted independent of the connection mode or any type of processing error, e.g., authentication failure.
- **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 for example no packets were received on the connection.

- **Interarrival jitter.** An estimate of the statistical variance of the RTP data packet interarrival time measured in milliseconds 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 [2]. The count includes packets received from different SSRC if the sender used several values. The value is zero if, e.g., no packets were received on the connection.
- **Average transmission delay.** An estimate of the network latency expressed in milliseconds. 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 [2].

In addition to the parameters above, an endpoint that has received one or more RTCP sender or receiver reports from its peer MUST return the following parameters:

- **Remote Packets Sent:** The number of packets that were sent on the connection from the perspective of the remote endpoint.
- **Remote Octets Sent:** The number of octets that were sent on the connection from the perspective of the remote endpoint.
- **Remote Packets Lost:** The number of packets that were not received on the connection, as deduced from gaps in the sequence number from the perspective of the remote endpoint.
- **Remote Jitter:** The average inter-packet arrival jitter, in milliseconds, expressed as an integer number from the perspective of the remote endpoint.

The **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. Note that absence of the **RequestedEvents** and **SignalRequests** parameters is interpreted as an empty list only if a **RequestIdentifier** parameter is included.

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.

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

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see Section 7.5) optionally followed by commentary.

7.3.6 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. A Call Agent that receives a DeleteConnection with a wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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 Section 7.6.

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 Section 7.5) optionally followed by commentary.

7.3.7 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. An embedded client that receives a **DeleteConnection** (Multiple Connections From the Call Agent) with the "any of" wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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 Section 7.1.1. 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 Section 7.5) optionally followed by commentary.

7.3.8 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 a PacketCable 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 this specification.

7.3.8.1 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]
  [, MaxMGCPDatagram]
  [, Capabilities] }
← AuditEndPoint(EndPointId
                  [, RequestedInfo] |
                  { [, SpecificEndPointID]
                    [, MaxEndPointIDs] })
```

The **EndPointId** identifies the endpoint that is being audited. The "any of" wildcard convention MUST NOT be used. An embedded client that receives an AuditEndPoint with the "any of" wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response.

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, ConnectionIdentifiers, DetectEvents, ObservedEvents, EventStates, VersionSupported, 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.

If an endpoint is queried about a parameter it does support, but has no value for, the endpoint MUST NOT generate an error; instead the parameter MUST be included in the response with an empty parameter value.

Only when successful, the response MUST, in turn, include information about each of the items for which auditing information was requested. Excluding parameters that are explicitly marked "optional", endpoints MUST support all of the following parameters:

- **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 signals²¹. 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.
- **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. Note that the MTA MAY include only the absolute domain name (including the hostname) of its NotifiedEntity if only the absolute domain name (including the hostname) was provided to it via the NotifiedEntity parameter of an NCS message or acknowledgement. The CMS SHOULD accept the value in this case.
- **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 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 Reason-Code 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 normal.
- **MaxMGCPDatagram:** The maximum size of an MGCP datagram in bytes supported by the endpoint (see Section 8.5.3). The value excludes any lower layer overhead. Support for this parameter is optional. The default maximum MGCP datagram size is assumed if a value is not returned.

²¹ Currently, there should be no pending brief signals.

- **Capabilities:** The capabilities for the endpoint similar to the LocalConnectionOptions parameter and including event packages and connection modes. If any unknown capabilities are reported, they MUST simply be ignored. 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 literal names defined in Table 4 the PacketCable Audio/Video Codecs Specification [11] MUST be used. Unknown compression algorithms SHOULD be ignored if they are received. 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 PacketCable 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: A parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see Section 7.5) 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.

7.3.8.2 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. An embedded client that receives a AuditConnection with a wildcard convention MUST return an error (the error returned SHOULD be error code 500 - the transaction could not be executed because the endpoint is unknown) in response. 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.
  
```

If an endpoint is queried about a connection parameter it does not support, the endpoint MUST NOT generate an error; instead the parameter MUST be omitted from the response.

If an endpoint is queried about a connection parameter it does support, but has no value for, the endpoint MUST NOT generate an error; instead the parameter MUST be included in the response with an empty parameter value. Only when successful, the response MUST, in turn, include information about each of the items for which auditing information was requested. Excluding parameters that are explicitly marked "optional", endpoints MUST support all of the following parameters:

- **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 most recent RemoteConnectionDescriptor that was supplied in a previous CreateConnection or ModifyConnection command to the gateway for this connection.

ReturnCode is a parameter returned by the gateway. It indicates the outcome of the command and consists of an integer number (see Section 7.5) 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).

7.3.9 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. A Call Agent that receives a Restart in Progress with the "any of" wildcard convention **MUST** return an error (the error returned **SHOULD** be error code 500 - the transaction could not be executed because the endpoint is unknown) in response. 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. At the expiry of the restart delay, the gateway should send a new RSIP message with a restart method of "forced". This will explicitly indicate to the Call Agent that the endpoints are now out of service.
- A "cancel-graceful" restart method indicates that a gateway is canceling a previously issued "graceful" restart method for the same endpoints. The endpoints remain in service. 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 endpoint will never go out of service as a result of this operation, and 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. However, an embedded client **MUST** send a "forced" RestartInProgress message when it is shutdown through the provisioning system. The Call Agent cannot rely on 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 Section 7.4.3.5 – 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 Section 7.4.3.6. The "restart delay" parameter **MUST NOT** be used with the "forced" and "cancel graceful" 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 Section 7.5) 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 Section 7.4.3.5 and 7.4.3.6). If a NotifiedEntity parameter was included in the response returned, it specifies a new "notified entity" for the endpoint(s) – this operation **SHOULD** only be done with the response error code 521 (endpoint redirected). Note that the behavior for returning a NotifiedEntity in the response is only defined for RestartInProgress responses and **SHOULD NOT** be done for responses to other commands. Any other behavior is undefined.

- 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 NotifiedEntity 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 "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 is terminated if a permanent (5xx) error code is returned, except for 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).

7.4 States, Failover and Race Conditions

In order to implement proper call signaling, 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.

7.4.1 Recaps and Highlights

As mentioned in Section 7.1.4, 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 section we recap and highlight the areas that are of special importance to reliability and failover 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 are 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 explicitly²², 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.
- 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 Agent's is not provided. We are relying strictly on the Call Agent's knowing what they are doing and communicating with each other (although AuditEndpoint can be used to learn about the current "notified entity").

7.4.2 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 Section 8.5), commands are repeated. MGCP Entities MUST keep in memory a list of the responses that they sent to recent transactions (i.e., a list of all the responses they sent over the last T_{hist} seconds), and a list of the transactions that have not yet finished executing. The default value for T_{hist} is 30 seconds.

The transaction identifiers of incoming commands are first compared to the transaction identifiers of the recent responses. If a match is found, the MGCP Entity 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 MGCP Entity does not execute the transaction; subsequent handling depends on the command in question. If it is a CreateConnection or ModifyConnection command, the MGCP Entity (gateway in this case) MUST send a provisional response. If it is any other command, it is simply ignored. In either case, a final 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.

²² This could for instance happen by specifying an empty NotifiedEntity parameter.

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". Max2 MUST be set to a larger value than Max1.

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: (A retransmission algorithm including these modifications is illustrated in Figure 4 below.)

- 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, then retransmissions MUST cease. If more than $2 * T_{thist}$ has elapsed, then 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. Also, receipt of explicit network notifications such as, e.g., ICMP network, host, protocol, or port unreachable SHOULD lead the gateway to try alternate addresses (with due consideration to possible security issues).
- 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 there still are no more interfaces to try, then retransmissions MUST cease. If more than $2 * T_{thist}$ has elapsed, then the endpoint becomes disconnected.
- Once an endpoint becomes disconnected, subsequent processing depends on whether the loss of association was detected by the gateway or Media Gateway Controller, as follows:
 - The gateway MUST initiate the "disconnected" procedure as specified in Section 7.4.3.6.
 - The Call Agent MUST NOT attempt to use the endpoint for any new calls until connectivity has been restored. Furthermore, the Call Agent MUST implement an algorithm to detect when connectivity with the endpoint is subsequently restored (e.g., on receiving a response to a periodic Audit Endpoint command). When connectivity to the endpoint is restored, and if no other condition exists which prevents the endpoint from supporting calls, the Call Agent MUST ensure that the endpoint can be used for new calls without requiring any manual intervention.

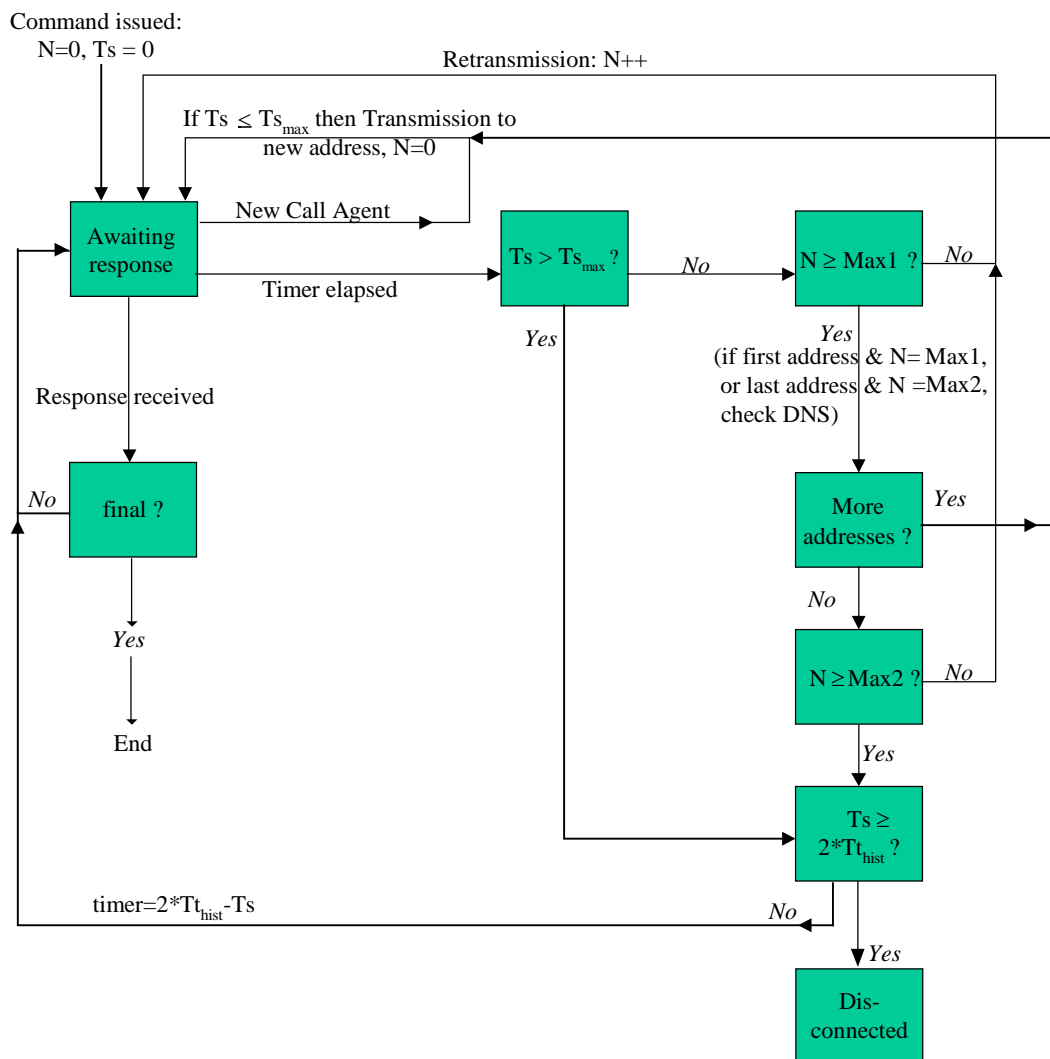


Figure 4 - Retransmission Algorithm

In order to adapt to network load automatically, MGCP specifies exponentially increasing timers (see Section 8.5.2). If the initial time-out is set to 200 milliseconds, the loss of a fifth retransmission will be detected after about 6 seconds. 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 seconds is probably acceptable.

It is however important that the maximum delay of retransmissions be bounded. Prior to any retransmission, it is checked that the time (T_s) elapsed since the sending of the initial datagram is no greater than $T_{s_{max}}$. If more than $T_{s_{max}}$ time has elapsed, retransmissions MUST cease. When $T_{s_{max}}$ has expired, or all of the retransmissions to all known IP addresses have been sent, there is a pause before declaring the endpoint disconnected. This pause represents a period of time where the only action is to wait for a response from any of the recent retransmissions. The quiescent period lasts for what remains of twice the life expectancy of the original transaction ($2 * T_{t_{hist}}$). This settling time allows all active transactions to complete or time out before the endpoint is declared disconnected. This helps ensure that each restart of the endpoint is from a clean and initial state. If more than $2 * T_{t_{hist}}$ time has elapsed, the endpoint becomes disconnected. The value $T_{s_{max}}$ is related to the $T_{t_{hist}}$ value: the $T_{t_{hist}}$ value MUST be greater than or equal to $T_{s_{max}}$ plus the maximum propagation delay in the network, $T_{p_{max}}$.

In other words, the following relation **MUST** be satisfied to prevent retransmitted commands from being executed more than once: $T_{hist} \geq T_{smax} + T_{pmax}$.

The default value for T_{smax} is 20 seconds. Thus, if the assumed maximum propagation delay is 10 seconds, then responses to old transactions must be kept for a period of at least 30 seconds. 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.

7.4.3 Race Conditions

In this section 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 behavior 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 reestablished, many gateways may decide to send RestartInProgress commands simultaneously, which could lead to very unstable operation if not carefully controlled.

7.4.3.1 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 Section 7.4.2. If the endpoint is or becomes disconnected (see Section 7.4.2) 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 Section 7.4.3.6 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 behavior 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 behaviors:

- 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,
- 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 reenter 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 NotificationRequest is here viewed as receiving a new NotificationRequest as well, except that the current list of ObservedEvents remains unmodified rather than being processed again.

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 the 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 event is reset to a null value. The subsequent behavior 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 only exception is the activation of an embedded Notification Request. In this case the observed event list remains unmodified rather than being processed again.

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 5 below illustrates the procedure specified above assuming all transactions execute successfully.

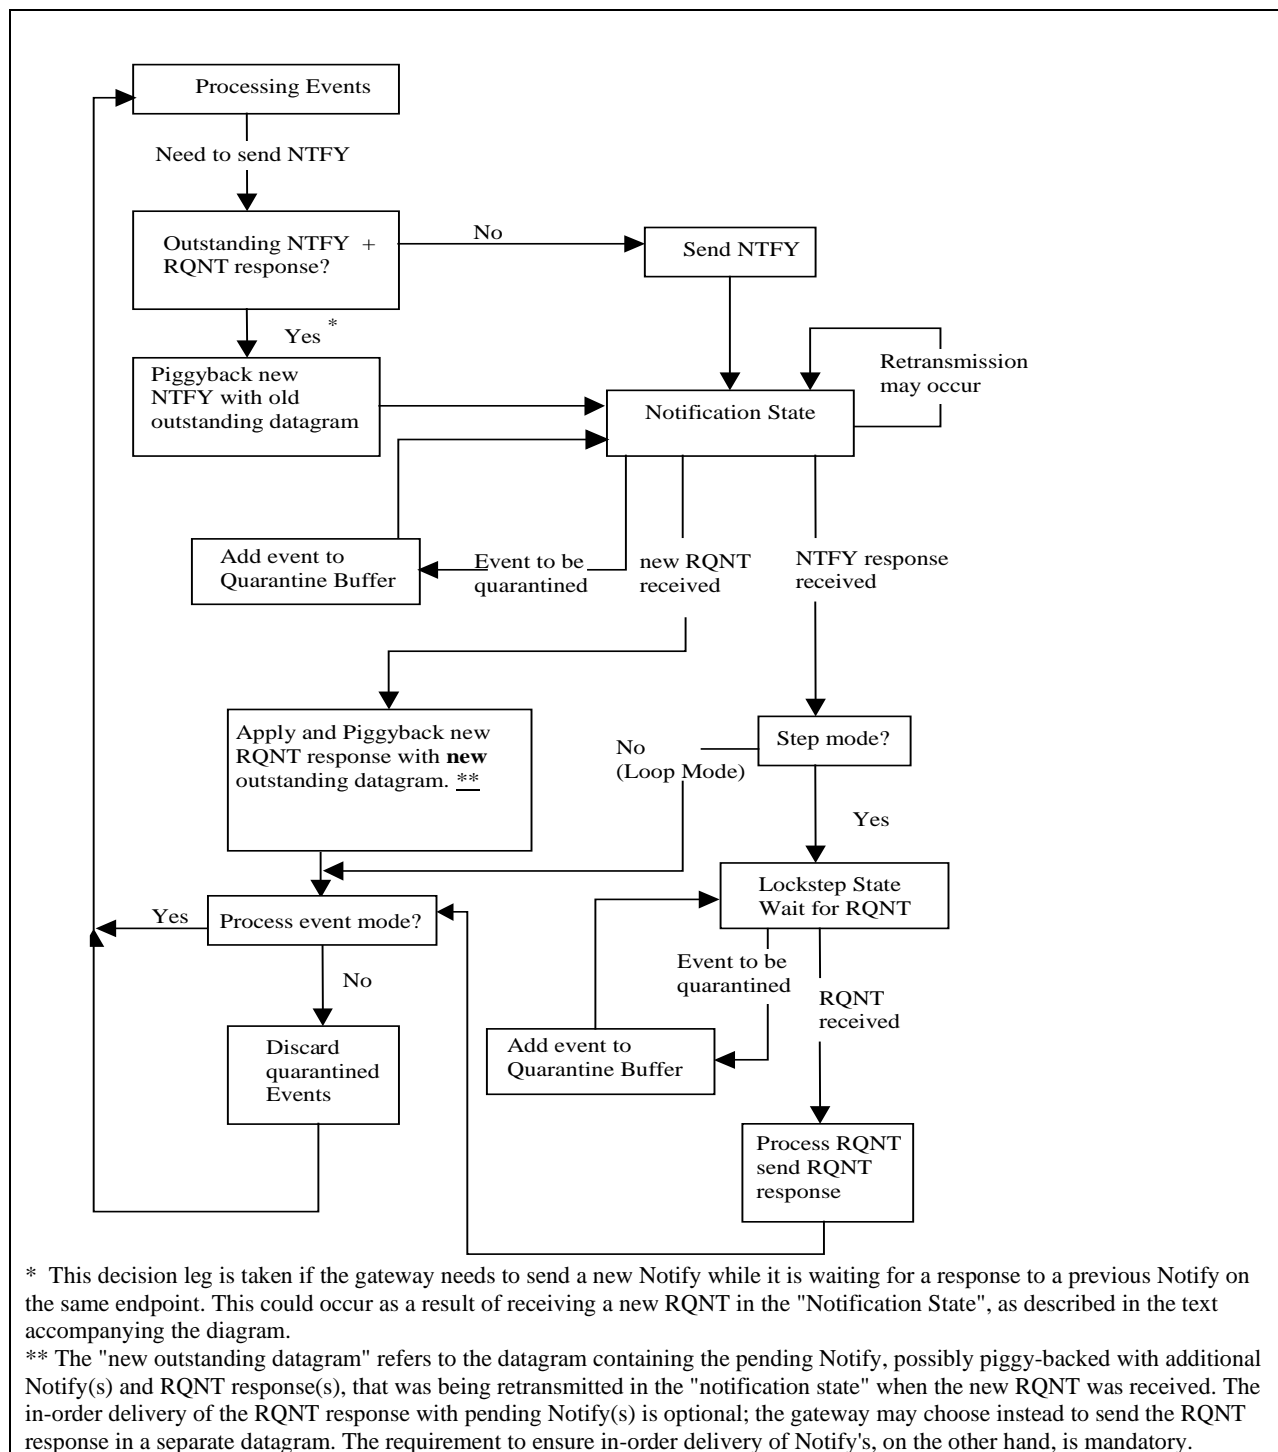


Figure 5 - Quarantine List Procedures

Call Agents SHOULD provide the response to a successful Notify message and the new NotificationRequest in the same datagram using the piggy-backing mechanism.²³

²³ Vendors who choose not to follow this recommendation should examine call agent failure scenarios carefully.

7.4.3.2 *Explicit Detection*

A key element of the state of several endpoints is the position of the hook. Although hook-state changing events are persistent in NCS, race conditions and state mismatch may still 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:

- If the gateway is requested to notify an "off hook" transition while the phone is already off hook (error code 401 – phone off hook),
- 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, whereas 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.

When the Call Agent receives an error response indicating that a **NotificationRequest** was unsuccessful, it **MUST** take action to ensure that any events quarantined by the endpoint are processed or discarded, and that the endpoint is returned to a normal operating mode where newly requested events are reported as they occur. For example, say the Call Agent receives an error code "401 - phone is already off hook" response to a **NotificationRequest** requesting detection of the "off-hook" signal. At this point the Call Agent should assume that the **NotificationRequest** had no effect on the endpoint, and that the endpoint is in the same state that it was prior receiving the command. If the endpoint was quarantining events in the "lockstep" state prior to receiving the **NotificationRequest**, then it will still be quarantining events in the "lockstep" state after sending the "401" error response. To ensure that the endpoint is not left in a state where it is permanently quarantining events, the Call Agent should send a new **NotificationRequest** with a different (possibly empty) set of requested events to transition the endpoint out of the "lockstep" state and into the normal mode where it can report new events.

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**.

7.4.3.3 *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.

7.4.3.4 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 below:

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 Section 7.4.3.5). 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 behavior **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.

To ensure consistent operation, embedded clients **SHOULD** behave as specified below when one or more of the above rules are not followed:

- Where a single outstanding command is expected (ModifyConnection, NotificationRequest), but the same command is received in a new transaction before the old finishes executing, the gateway **SHOULD** fail the previous command. This includes the case where one or more of the commands were encapsulated. The use of error code 407 (transaction aborted) is **RECOMMENDED**.
- If a ModifyConnection command is received for a pending CreateConnection command, the ModifyConnection command **SHOULD** simply be rejected. The use of error code 400 (transient error) is **RECOMMENDED**. Note that this situation constitutes a Call Agent programming error.

Note, that where reception of a new command leads to aborting an old command, the old command **SHOULD** be aborted regardless of whether the new command succeeds or not. For example, if a ModifyConnection command is aborted by a DeleteConnection command which itself fails due to an encapsulated NotificationRequest, the ModifyConnection command is still aborted.

7.4.3.5 *Fighting the Restart Avalanche*

Let's 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 behavior **MUST** be followed:

1. When a gateway is powered on, or when a gateway's endpoints are put into service, the gateway initiates a restart timer to a random value, uniformly distributed between 0 and a provisionable maximum waiting delay (MWD), e.g., 360 seconds (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 RestartInProgress 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 behavior 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, then the endpoint is redirected, and the restart procedure **MUST** be initiated again (as a new transaction). The 521 response should have included a NotifiedEntity which then is the "notified entity" towards which the restart is initiated.
- If the error is any other permanent error (5xx), 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 endpoint enter the "disconnected" state while carrying out the restart procedure, then the disconnected procedure specified in Section 7.4.3.6 **MUST** be carried out, and a "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 minutes. 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 minutes on average, or, to put it otherwise, about one transaction per endpoint every 5 to 6 minutes on average. This suggests that a reasonable value of MWD for a residential gateway would be 10 to 12 minutes. In the absence of explicit configuration, embedded clients **MUST** use a default value of 600 seconds for MWD.

7.4.3.6 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 Section 7.4.2. 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 seconds. 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 ($T_{d_{min}}$) 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, a new value for the "disconnected" timer is selected. The timer value **MUST** be selected from the range defined by 1.5 times the last timer value and double the last timer value, and **MAY** be randomly generated. In either case, the new timer value is subject to a provisionable "disconnected" maximum waiting delay ($T_{d_{max}}$), e.g., 600 seconds, 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 Section 7.4.2.) The endpoint **MUST** take full advantage of piggy-backing in achieving this.

On receiving a RestartInProgress message containing a restart method of "disconnected" the Call Agent **MUST** take action to ensure that any events quarantined by the endpoint are processed or discarded, and that the endpoint is returned to a normal operating mode where newly requested events are reported as they occur. The Call Agent **SHOULD** send a NotificationRequest containing a Quarantine-Handling parameter set to "discard" in this case.

The Call Agent may also decide to perform one or more of the following; audit the endpoint, clear all connections for the endpoint, or send a NotificationRequest asking the endpoint to process quarantined events (see Section 7.4.3.7).

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.

Note also that a disconnected endpoint does not mean that the endpoint is in an "out-of-service" state. Disconnected is not a state of the endpoint's service availability, but rather an indication of the endpoints inability to communicate to the call agent.

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 (Section 7.4.3.5). If the "disconnected" procedure is to be initiated again following an error response, the rate-limiting timer considerations specified above still apply. 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 "disconnected" 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 endpoints MUST NOT blindly drop new commands to be sent for a period of $T_{s_{max}}$ seconds after they receive a non-audit command. One way of satisfying 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 for $T_{s_{max}}$ 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 $T_{s_{max}}$ seconds, but a "disconnected" procedure is initiated within $T_{s_{max}}$ seconds, the endpoint MAY piggy-back the buffered command(s) with that RestartInProgress. Note, that once a command has been sent, regardless of whether it was buffered initially, or piggy-backed earlier, retransmission of that command MUST cease $T_{s_{max}}$ seconds after the initial send as described in Section 7.4.2. This specification purposely does not specify any additional behavior 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 $T_{d_{init}}$ is 15 seconds, the default value for $T_{d_{min}}$, is 15 seconds, and the default value for $T_{d_{max}}$ is 600 seconds.

7.4.3.7 Call Agent Handling of Disconnected Endpoints

When an endpoint is in the "disconnected" state it may accumulate a large number of events in the quarantine buffer. Also, a "disconnected" endpoint may autonomously delete established connections (say the gateway reboots). Therefore, when connectivity between a "disconnected" endpoint and its Call Agent is subsequently restored, the Call Agent MUST be prepared to deal with the following issues:

- The large number of Notify messages that may be generated by the endpoint if all events on the quarantine list are processed,
- The reception of old/stale events reported by the endpoint that no longer have any relevance. The quarantine buffer is a First-In-First-Out (FIFO) queue, where the oldest events are processed first and notified (if requested) to the Call Agent. The action taken by the Call Agent on receiving an old event may not be meaningful if the old event has been superseded by newer events (for example, a "digits" event would no longer be relevant if the endpoint subsequently went on-hook),
- Connection state mismatch between the Call Agent and endpoint, where the Call Agent thinks the endpoint has one or more connections, but the endpoint in fact has no connections.

Call Agents are free to use any mechanism supported by the protocol to deal with the above issues. One way of achieving this is to do the following:

1. We define a new boolean variable called "disconnect-event-sync", which is maintained by the Call Agent for each of its endpoints. When set to true, this variable indicates that connectivity with a "disconnected" endpoint has recently been restored, but that event/signal synchronization has not yet been achieved. (Note, this variable

is being introduced in order to describe Call Agent behavior, and is not intended to imply any particular implementation. The variable is not externally visible.)

- a. As soon as the Call Agent learns that an endpoint is disconnected, it sets "disconnect-event-sync" to true. The "disconnected" procedure ensures, that the Call Agent will learn about the endpoint being disconnected by receiving a "disconnected" RestartInProgress message. When an endpoint receives a positive acknowledgement to the "disconnected" RSIP, it completes the "disconnected" procedure. At this point, the endpoint could immediately generate a Notify for a two reasons; to send a Notify command that was buffered while the endpoint was disconnected, or, if the endpoint was in the "notification" state and "loop" mode while disconnected, to report a Notify-triggering event on the quarantine list.

- If the endpoint operates in "step" mode, responding to the Notify will in itself not enable any further Notify messages to be generated (an additional NotificationRequest would be needed for that).
- However, if the endpoint operates in "loop" mode, then a response to the Notify will enable further Notify messages to be generated. As discussed above, this is sometimes undesirable as the events being reported may be old and a lot of events may have been quarantined, which in turn will result in a large number of Notify messages and subsequent NotificationRequest messages based on stale information.

- b. As long as an endpoint has "disconnect-event-sync" set to true, the Call Agent should take action to ensure that the potentially large number of events on the quarantine list are either discarded, or processed in a controlled and orderly manner. This can be achieved in a number of ways:

- The Call Agent can send a single NotificationRequest specifying that all quarantined events are to be discarded. On receiving a positive acknowledgment to this command, or on receiving a Notify with the same RequestIdentifier, the Call Agent should set the "disconnect-event-sync" indicator to false, at which point regular event processing for the endpoint is resumed. The downside to this approach is that it will clear any events that have been accumulated, regardless of the number of events accumulated. In some cases, this may result in unnecessary service interruption. In order to address this, protocol extensions will be necessary.
- The Call Agent can send a NotificationRequest specifying that quarantined events are to be processed. If the endpoint is operating in "step" mode then it will report a single notify-triggering event for each NotificationRequest received, whereas in the "loop" mode it can report multiple events under a single NotificationRequest.

Since the information conveyed by the notified events may no longer be relevant, the Call Agent should not blindly process these events (for example, on being notified of an "off-hook" event, the Call Agent should not automatically send a NotificationRequest to apply dial-tone and request notification of "on-hook"). Instead, the Call Agent must synchronize its internal state data with the actual state of the endpoint. Since an endpoint performs signal/event glare processing against the current hook-switch state of the line, the Call Agent can discover the current hook-switch state based on the response to a NotificationRequest. For example, a "402 - phone already on hook" response to a NotificationRequest asking for detection of "on-hook" implies that the line is currently on-hook. The Call Agent may choose to ignore notified events that are deemed irrelevant based on the current hook-switch state (e.g., ignore "digits" event if the line is currently "on-hook").

Once all events on the quarantine list have been processed, the Call Agent should set the "disconnect-event-sync" indicator to false. The Call Agent can safely assume that all quarantined events have been processed if a delay of duration T_{thist} has expired since the Call Agent last prompted the endpoint to process the next event (i.e., if a delay of duration T_{thist} has expired since the Call Agent sent the response to the previous Notify in the "loop" mode, or since the Call Agent received the last positive response to NotificationRequest in the "step" mode).

2. When endpoints become disconnected, connections created on that endpoint should be unaffected. However, it is always possible, that a connection can no longer be sustained by the endpoint and hence that it is deleted; this will result in a DeleteConnection command sent to the Call Agent. When the endpoint is disconnected, such a command may never make it to the Call Agent, in which case it will not learn about the deleted connection.

Consequently, whenever a Call Agent learns that an endpoint is disconnected, it should audit the endpoint for the list of connections present on it.

7.5 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 five value ranges have been defined:

- value 000 indicates a response acknowledgement²⁴,
- 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.

The values that have been defined are listed in Table 4 below:

Table 4 - Return Code Definitions

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.
407	Transaction aborted. The transaction was aborted by some external action, e.g., a ModifyConnection command aborted by a DeleteConnection command.
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.
503	"All of" wildcard not fully supported. The transaction contained an "all of" wildcard, however the gateway does not fully support these. Note that this is currently only permissible for non-empty NotificationRequests.
505	Unsupported RemoteConnectionDescriptor. This SHOULD be used when one or more mandatory parameters or values in the RemoteConnectionDescriptor is not supported.
506	Unable to satisfy both LocalConnectionOptions and RemoteConnectionDescriptor. This SHOULD be used when the LocalConnectionOptions and RemoteConnectionDescriptor contain one or more mandatory parameters or values that conflict with each other and/or cannot be supported at the same time (except for codec negotiation failure – see error code 534).
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.

²⁴ Response acknowledgement is used for provisional responses (see Section 8.8)

Code	Meaning
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.
538	Event/signal parameter error (e.g., missing, erroneous, unsupported, unknown, etc.).

7.6 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. Defined values of the reason code are listed in Table 5 below:

Table 5 - Reason Code Definitions

Code	Meaning
000	Endpoint state is normal. (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.

7.7 Use of Local Connection Options and Connection Descriptors

The normal sequence in setting up a bi-directional connection involves at least three 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 toward 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 ModifyConnection 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 honor the media parameters until either the connection is deleted, or the gateway issues a new LocalConnectionDescriptor. In addition to the media parameters assigned to the connection, the gateway can also advertise additional supported capabilities in the LocalConnectionDescriptor. Note that such capabilities MUST be provided outside the "m=" line in the SDP. The gateway is free to advertise all of its supported capabilities independent of the LCO or RCD parameters received from the Call Agent, and independent of the media parameters associated with the connection.

Codec and packetization period selection must only be performed, as described in this section, if either:

- a) The gateway receives a CRCX, or
- b) The gateway receives a MDCX and any of the following parameters are present:
 - encoding method (a: in LocalConnectionOptions)
 - packetization period (p: in LocalConnectionOptions)
 - multiple packetization period (mp: in LocalConnectionOptions)
 - RemoteConnectionDescriptor

Furthermore, this codec and packetization period selection process must only use the information present in the connection request and not retain any of the values that may have been received in previous connection requests. For example, if a gateway received a MDCX with all necessary LCO parameters but was missing a RemoteConnectionDescriptor, it will negotiate as if no RemoteConnectionDescriptor had ever been received for that connection. As well, if all of the above parameters are omitted in a MDCX command the existing negotiated codecs and packetization periods will remain intact.

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

- A list of codecs and packetization periods allowed by the LocalConnectionOptions. A codec is allowed by the LocalConnectionOptions if it satisfies the constraints specified by the encoding method, packetization period and multiple packetization periods fields. If one or more of these fields are omitted, the omitted fields do not impose any constraints on the allowed codecs.
- A list of codecs and packetization periods in the RemoteConnectionDescriptor.

- An internal list of codecs and packetization periods that the gateway can support for the connection. A gateway may support one or more codecs and packetization periods 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/packetization periods is formed by taking the intersection of the internal list of codecs/packetization periods and codecs/packetization periods allowed by the LocalConnectionOptions. If LocalConnectionOptions was not provided, the approved list of codecs/packetization periods thus contains the internal list. If the LocalConnectionOptions was provided but the codecs parameter was omitted, the LocalConnectionOptions implicitly allows all codecs in the internal list, provided they are not incompatible with any packetization period(s) specified. Similarly, if the LocalConnectionOptions was provided but the packetization period(s) was omitted, the LocalConnectionOptions implicitly contains the set of packetization periods supported by the internal list.
2. If the approved list of codecs/packetization periods 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/packetization periods is formed by taking the intersection of the approved list of codecs/packetization periods and codecs/packetization periods allowed by the RemoteConnectionDescriptor. If a RemoteConnectionDescriptor was not provided, the negotiated list of codecs/packetization periods thus contains the approved list of codecs/packetization periods. If the RemoteConnectionDescriptor does not contain any media stream lines, a codec negotiation failure has occurred, and an error response is generated (error code 534 – codec negotiation failure – is recommended). If the RemoteConnectionDescriptor contains multiple media streams, the MTA SHOULD only accept one of these and reject the others by setting their port to zero in the LocalConnectionDescriptor. If the RemoteConnectionDescriptor was provided but the packetization period(s) was omitted, the negotiated list of packetization periods contains the set of packetization periods from the approved list. The MTA MUST choose reasonable defaults as referenced or defined in the PacketCable 1.5 Audio/Video Codecs Specification [11] if the packetization period is explicitly omitted from both the LocalConnectionOptions and the RemoteConnectionDescriptor.
4. If the negotiated list of codecs/packetization periods 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/packetization periods 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. It should be noted, that the above procedure negotiates both encoding methods and packetization periods as opposed to just encoding methods. This is done to enable consistent local and far-end QoS operation in the segmented QoS model used in PacketCable.

In the case that a gateway supports more than one codec per endpoint, there are two options the gateway can use in deciding how many codecs it wants to support for that connection:

1. Gateway supports multiple codecs and can switch between different codecs in real-time. The gateway returns all negotiated codecs in the SDP media stream line and reserves the Least-Upper-Bound (LUB) as per the DQoS specification [12]. The LUB is reserved to guarantee that a switch to any of these codecs will succeed. Multiple codecs in the m= line means the device must be ready to receive media packets from any of the negotiated codecs. As well, the gateway may send media packets from any of the negotiated codecs and switch between them as required.
2. Gateway supports one or more codecs but cannot switch between different codecs in real-time. The gateway therefore negotiates and returns only one codec in the SDP media stream line, (optionally, gateway also puts additional supported codecs in the SDP 'X-pc-codecs' attribute) and reserves the bandwidth for the single negotiated codec in the media stream line as per the DQoS specification. With this method, a codec change must be initiated by the CMS in order to change codecs at which time the resulting change in bandwidth is re-established as per the DQoS specification.

7.7.1 RFC 2833 Negotiation

The internal list of supported codecs **MUST** include the telephone-event codec with events 0-15. This will ensure that RFC 2833 DTMF relay will be used for the connection when authorized by the LCO (via inclusion or an empty a: parameter) and allowed by the RCD.

An example of an LCO that authorizes RFC 2833 DTMF relay is:

```
L: a:PCMU;PCMA;telephone-event, mp:10;20;-
```

The telephone-event codec **MUST NOT** be the only codec supplied in the LCO from the CMS. If the MTA receives an LCO containing only the telephone-event codec, then it **MUST** return error code 524 – internal inconsistency in LocalConnectionOptions. If the approved list of codecs, as described in Section 7.7, contains only the telephone-event codec, the endpoint **MUST** return error code 534 – codec negotiation failure. Likewise, if the negotiated list of codecs, as described in Section 7.7, contains only the telephone-event codec the endpoint **MUST** return error code 534.

If the packetization period LCO parameter is used for the connection, the endpoint **MUST** use that packetization rate for the DTMF relay packets. If the multiple packetization period parameter is used in the LCO, the CMS **MUST** use a hyphen for the packetization rate for the telephone-event codec. If the endpoint receives an LCO with the multiple packetization period parameter with a packetization rate for telephone-event codec not set to a hyphen, the endpoint **MUST** return error code 524 – Inconsistent LCO. When an endpoint returns an LCD which includes the ability to receive the telephone-event codec, it **MUST** use a hyphen as the packetization rate in the mptime SDP attribute.

An example of an LCD which advertises support of RFC 2833 DTMF digits is:

```
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 0 8 105
a=mptime:10 20 -
a=rtpmap:105 telephone-event/8000/1
```

For more information on the use of RFC 2833 DTMF Digits refer to PacketCable Audio/Video Codec Specification [11].

7.7.2 Selecting Codec Packetization Interval for T.38

The packetization interval for T.38 is selected using the same procedure that is used for audio codecs. See section 7.7 above.

Special DQoS considerations are required for T.38. When T.38 is allowed and DQoS is to be utilized the CMS **SHOULD** consider the bandwidth requirements for T.38 at a particular packetization period in its LUB calculation in order to authorize enough bandwidth in case of a later switch to T.38. This packetization rate is known by the CMS – either by default or via configuration. If the CMS later instructs the endpoint to transition to T.38 via a MDCX command using the same gate, the packetization period used for T.38 – either specified explicitly in the Local Connection Options or assumed by default - **MUST** be within the DQoS resources authorized previously. If the current gate cannot satisfy a request for bandwidth by the MTA, the CMS **MUST** establish a new gate capable of satisfying such a request and relay the new GateID to the MTA in the MDCX.

The endpoint, however, **MUST NOT** include T.38 in its LUB calculation upon receipt of a CRCX or MDCX command when DQoS is requested unless "image/t38" is selected through the codec negotiation procedure. If only "image/t38" is specified as the encoding method in the Local Connection Options of either a CRCX or MDCX command, then the endpoint **MUST** consider T.38 at the specified packetization period only in its LUB calculation. Even though the far end may send T.38 packets at a higher packetization period than that authorized by the local Call Agent as specified by the PacketCable Audio/Video Codecs Specification [11] the endpoint **MUST NOT** consider these higher packetization periods in its calculation for bandwidth utilization in the downstream. There will already be enough bandwidth to support the potentially higher packetization period.

7.7.3 Remote IP and Port Negotiation

The remote IP and port are provided via the Remote Connection Descriptor. Once obtained in a successful connection handling command (for example: a Modify Connection command), the endpoint **MUST** continue to use them until either a new Remote Connection Descriptor is provided which specifies a new remote IP or port, or the connection is deleted. Note that the receipt of a Modify Connection command without a Remote Connection Descriptor which may kick off Codec negotiation, does not invalidate the current remote IP and port, even if the media is modified from audio to image or vice versa.

7.7.4 LCO Use with FXR Package

The following fxr/fx LCO formats **MAY** be used if FXR is used without VBD package (see Appendix A.5):

1. fxr/fx:gw

This results in image/t38 via udptl t38 not being used on the call. Other ways for the CMS to request the same behavior if the VBD package is not being used are:

fxr/fx:off

no fxr/fx in the LCO (this provides the default of gw)
or any fxr/fx format where off is first followed by one or more other methods.

2. fxr/fx:t38;off

image/t38 via udptl t38 or tcp t38 is used if the rcd has 3407 indication of image/t38 capability, otherwise the fax is carried via audio codecs.

Other ways for the CMS to request the same behavior when the VBD/GPMD package is not used are:

fxr/fx:t38;gw

fxr/fx:gw;t38

3. fxr/fx:t38

image/t38 via udptl t38 or tcp t38 is used if the rcd has 3407 indication of image/t38 capability and fax is detected. If rcd does not have a 3407 indication of image/t38, the connection request is to be nacked.

4. fxr/fx:t38-loose

Other ways for the CMS to request the same behavior when the vbd package is not used are:

fxr/fx:gw;t38-loose

fxr/fx:t38-loose;off

fxr/fx:t38-loose;gw

fxr/fx:t38-loose;t38;off

fxr/fx:t38-loose;t38;gw;off

fxr/fx:t38-loose;gw;t38;off

etc.

However, any use of fxr/fx: with two or more methods where t38-loose if the first method listed is **NOT RECOMMENDED**.

7.7.5 LCO with FXR and VBD Being Used by CMS at Same Time

If FXR and VBD packages are used simultaneously, the LCO format used depends on whether T.38 or V.152 is the preferred method of handling fax.

1. If t38-loose is the preferred method of handling fax, then the LCO MUST include fxr/fx as formatted in the following example:
 L:a:PCMU;red;PCMU, fxr/fx:t38-loose, gpmd/gpmd:"PCMU:2 vbd=yes", ftmp:"RED
 PCMU:2/PCMU:2".
2. If t38-strict is preferred for fax and V.152 is to be used only if t38-strict cannot be used, the LCO MUST include fxr/fx as formatted in the following example:
 L:a:PCMU;red;PCMU, fxr/fx:t38;gw, gpmd/gpmd:"PCMU:2 vbd=yes", ftmp:"RED
 PCMU:2/PCMU:2".
3. If V.152 is preferred over T.38-strict, then the LCO MUST include fxr/fx as formatted in the following example:
 L:a:PCMU;red;PCMU, fxr/fx:gw;t38, gpmd/gpmd:"PCMU:2 vbd=yes", ftmp:"RED
 PCMU:2/PCMU:2"
 Other fxr/fx formats are NOT PERMITTED when both FXR and VBD packages are used at the same time.
4. If T.38 is not to be used AND the VBD package is being used, the CMS MUST either send fxr/fx:gw or simply omit fxr/fx from the LCO. The CMS MUST NOT include fxr/fx:off in the LCO when the VBD package is used - this is to ensure that the nopfax(start) event is not sent if V.152 is used for fax.
5. If V.152 is preferred over T.38-loose, the LCO MUST include fxr/fx as formatted in the following example:
 L:a:PCMU;red;PCMU, fxr/fx:gw;t38-loose, gpmd/gpmd:"PCMU:2 vbd=yes", ftmp:"RED
 PCMU:2/PCMU:2".

7.7.6 LCO format for V.152 and RFC 2198 Redundancy use

The CMS will specify in the LCO the payload to be used as V.152, the level of redundancy to be used with V.152 and the preference of the use of redundancy vs. non-redundant use of V.152.

The CMS SHOULD map V.152 to PCMU. If the CMS wishes to authorize the use of PCMU as voice and as V.152, it MUST map the second instance of PCMU in the a: section to V.152 as follows:

L:a:PCMU;red;PCMU, gpmd/gpmd:"PCMU:2 vbd=yes", ftmp:"RED PCMU:2/PCMU:2".

If Redundancy is specified by the CMS, it MUST be specified in the a: section ahead of the codec entry that is mapped to V.152.

The CMS MUST be capable of indicating a level of redundancy of "one" in the LCO. It does this by specify ftmp:"RED formatx/formatx".

Gateways that receive an LCO with Redundancy level greater than one MUST not nack the connection request but MAY reply with an lcd with redundancy level of one or without redundancy.

8 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 Section 7.3.6. 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.

8.1 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.

8.2 Command Header

The command header is composed of:

- A command line identifying the requested action or verb, the transaction identifier, the endpoint toward 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.

8.2.1 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.

8.2.1.1 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 6 below:

Table 6 - Requested Verb Codings

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).

8.2.1.2 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 Section 8.5). 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 this specification though.

Transaction identifiers for all commands sent from a given embedded client **MUST** be unique for the maximum lifetime of the transactions (see Section 8.5) 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. Transaction identifiers should not use any leading zeroes. Equality is based on numerical value and leading zeroes are ignored. An MGCP entity **MUST NOT** reuse a transaction identifier more quickly than three minutes after completion of the previous command in which the identifier was used.

8.2.1.3 Endpoint, Call Agent and NotifiedEntity Name Coding

The endpoint names and Call Agent names are encoded as e-mail addresses, as defined in [14]. 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 shown in Table 7 below:

Table 7 - Example Name Coding

aaln/1@ncs2.whatever.net	Analog 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 Call Agent port 2727 will be used. Additional detail on endpoint names can be found in Section 7.1.1.

8.2.1.4 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 this specification is 1.0. The protocol version for this specification **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).

8.2.2 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 8 below:

Table 8 - Parameter Definitions

Parameter name	Code	Parameter value
ResponseAck ²⁵	K	See description.
CallId	C	Hexadecimal string, MUST NOT exceed 32 characters. Call Identifiers are compared as strings rather than numerical values.
ConnectionId	I	Hexadecimal string, MUST NOT exceed 32 characters. Connection Identifiers are compared as strings rather than numerical values.
NotifiedEntity	N	An identifier, in RFC 2821 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 RFC 2821 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.
MaxMGCPDatagram	MD	See description.

²⁵ The ResponseAck parameter was not shown in Section 7.3 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 9 provides the association between parameters and commands. The letter M stands for mandatory, O for optional, and F for forbidden.

Table 9 - Association of Parameters with Commands

Parameter name	CRCX	MDCX	DLCX	RQNT	NTFY	AUEP	AUCX	RSIP
ResponseAck ²⁵	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*	O*	O*	O*	F	F	F	F
SignalRequests	O*	O*	O*	O*	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
MaxMGCPDatagram	F	F	F	F	F	F	F	F
RemoteConnection Descriptor	O	O	F	F	F	F	F	F

- * 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, omission of these two parameters when the command includes a RequestIdentifier means the corresponding lists will be considered empty.

Embedded clients and Call Agents SHOULD always provide mandatory parameters before optional ones, however embedded clients MUST NOT fail if this recommendation 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-ABC*.

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.

8.2.2.1 Response Acknowledgement

The response acknowledgement parameter²⁵ is used to support the three-way handshake described in Section 8.7. It contains a comma separated list of "confirmed transaction-id ranges".

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

8.2.2.2 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. RequestIdentifiers are compared as strings rather than numerical values. The string "0" is reserved for reporting of persistent events in the case where no NotificationRequest has been received yet (see Section 7.3.2).

8.2.2.3 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.
- The multiple packetization period in milliseconds for each codec in the encoding method LCO, encoded as the keyword "mp" followed by a colon and a list of decimal numbers or hyphens, with one entry for each entry in the Encoding Method Field. Each packetization period value is separated from its successor by a single semicolon. The first entry in the list MUST be a decimal number. Subsequent entries in the list MUST be either a decimal number or a hyphen.
- The literal name of the compression algorithm as specified in Table 4 of the Packet Cable Audio/Video Codecs Specification [11], encoded as the keyword "a" followed by a colon and a character string. If the Call Agent specifies a list of values, these values will be separated by a semicolon. For RTP, audio codecs MUST be specified by using encoding names defined in the RTP AV Profile [3], encoding names registered with the IANA, or encoding names referenced or defined in Table 4 of the Packet Cable Audio/Video Codecs Specification [11]. Non-audio media registered as a MIME type MUST use the "<MIME type>/<MIME subtype>" form, as in "image/t38". It is RECOMMENDED that other well-known variants of the literal codec names be supported as well.
- 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 shown in Table 10 below:

Table 10 - DQoS Resource Reservation Parameter Values

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 MUST be separated by a single 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 MUST be separated by a single 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 [18].

Whitespaces MUST NOT be sent within a ciphersuite, or between adjacent ciphersuites when multiple ciphersuites are provided. The following example illustrates the formatting of a ciphersuite and a ciphersuite list:

```
sc-rtp:62/51;64/51;60/50
```

The actual list of PacketCable supported ciphersuites is provided in the PacketCable Security Specification [13].

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:G729, 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).

8.2.2.4 Capabilities

Capabilities inform the Call Agent about the endpoint's 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 milliseconds, 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. The literal names defined in Table 4 of the PacketCable Audio/Video Codecs Specification [11] **MUST** be used. A list of values may be specified in which case the values will be separated by a semicolon.
- The bandwidth in kilobits per second (1000 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.²⁶
- 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 Section 8.2.2.7.
- The keyword "dq-gi" if Dynamic Quality of Service is supported.
- The keyword "sc-rtp" 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;p:10-30, e:on, s:off, v:L;S,
m:sendonly;recvonly;sendrecv;inactive

²⁶ Silence suppression on is equivalent to VAD enable and silence suppression off is equivalent to VAD disable.

```

A: a:G729; p:10-20, e:on, s:off, v:L;S,
  m:sendonly;recvonly;sendrecv;inactive
A: a:G729; p:30-90, e:on, s:on, v:L;S,
  m:sendonly;recvonly;sendrecv;inactive;confrnce,
  dq-gi, sc-rtp: 64/51;60/51, sc-rtcp: 71/81

```

Note that the codecs and security algorithms are merely examples – separate PacketCable specifications detail the actual codecs and algorithms supported, as well as the encoding used (see [11] and [13]). 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 this document.

8.2.2.5 Connection Parameters

Connection parameters are encoded as a string of type and value pairs, where the type is one of the codes in Table 11 below, and the value is a decimal integer. Types are separated from values by an "=" sign. Parameters are separated from each other by a comma.

Table 11 - Connection Parameters

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 milliseconds, expressed as an integer number.
Latency	LA	Average latency, in milliseconds, expressed as an integer number.
Remote Packets sent	PC/RPS	The number of packets that were sent on the connection from the perspective of the remote endpoint.
Remote Octets sent	PC/ROS	The number of octets that were sent on the connection from the perspective of the remote endpoint.
Remote Packets lost	PC/RPL	The number of packets that were not received on the connection, as deduced from gaps in the sequence number from the perspective of the remote endpoint.
Remote Jitter	PC/RJI	The average inter-packet arrival jitter, in milliseconds, expressed as an integer number from the perspective of the remote endpoint.

Extension connection parameter names are composed of the string "X-" followed by a two or three letter extension parameter name. Call Agents that receive unrecognized extensions MUST silently ignore these extensions. If an endpoint receives RTCP packets with these statistics, it MUST return the Remote parameters (Rxx above) in the response to the Delete-Connection and Audit-Connection commands.

An example of a connection parameter encoding is:

```

P: PS=1245, OS=62345, PR=0, OR=0, PL=0, JI=0, LA=48, PC/RPS=0, PC/ROS=0, PC/RPL=0,
  PC/RJI=0

```

8.2.2.6 Reason Codes

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

```

E: 900 Endpoint malfunctioning

```

A list of reason-codes can be found in Section 7.6.

8.2.2.7 Connection Mode

The connection mode describes the connection's operation mode. The possible values are shown in Table 12 below:

Table 12 - Connection Mode

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: replcate	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.

8.2.2.8 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.

Table 13 provides examples of event names:

Table 13 - Event Name Examples

L/hu	On-hook transition, in the line package.
L/0	Digit 0 in the Line package.
Hf	Flash-hook, assuming that the line package is the default package for the endpoint.
L/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): Table 14 provides examples of valid range and wildcard notation.

Table 14 - Event Range and Wildcard Notation

L/[0-9]	Digits 0 to 9 in the Line package.
L/X	Digits 0 to 9 in the Line package.
[0-9*#A-D]	All digits and letters in the Line package (default for endpoint).
L/all	All events in the 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. Table 15 provides examples of valid use of the star and dollar sign notations:

Table 15 - "All" and "Current" Connection Notation

L/ma@*	The RTP media start event on all connections for the endpoint.
L/rt@\$	Ringback on the current connection.

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

8.2.2.9 RequestedEvents

The RequestedEvents parameter provides the list of events that have been requested. The currently defined event codes are described in Appendix 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 Section 7.3.1 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 shown in Table 16 below:

Table 16 - Requested Events Actions

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, "ft" and "ft(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:

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:

R: hd(A, E(S(dl), R(oc(N), [0-9#T](D)), D((1xxxxxxxxxx|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:

R: hf(A, C(M(inactive(X43DC)), M(sendrecv(\$))), oc(N), 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".

8.2.2.10 SignalRequests

The SignalRequests parameter provides the name of the signals that have been requested. The currently defined signals can be found in Appendix A. A given signal can only appear once in the list, and all signals will, by definition, be applied at the same time. The MTA MUST support, at a minimum, a single signal on each endpoint and simultaneously support the generation of one signal on each connection for a given endpoint. Specific packages MAY define requirements beyond these minimum capabilities. For signal combinations beyond this minimum requirement that the MTA does not support, it SHOULD return error code 502.

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 Section 7.3.1):

- On/Off (OO),
- Time-out (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 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 for 6 seconds:

```
rg(to=6000)
rg(to(6000))
```

Individual signals may define additional signal parameters.

The signal parameters will be enclosed within parenthesis, as in (assuming "Line" is the default package):

```
S: ci(10/14/17/26, "555 1212", CableLabs)
```

When several signals are requested, their codes are separated by a comma, as in:

```
S: rg, rt@FDE234C8
```

8.2.2.11 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 are:

```
O: hu
O: ma@A43B81
O: 8,2,9,5,5,5,5,T
```

O: hf,hf,hu

O: 8,2,9,5,mt,5,5,5,T

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 Section 7.4.3.1 for details.

8.2.2.12 RequestedInfo

The RequestedInfo parameter contains a comma separated list of parameter codes, as defined in the "Parameter lines" section – Section 7.3.8 lists the parameters that can be audited. The values listed in Table 17 below are supported as well:

Table 17 - RequestedInfo Parameter Values

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

8.2.2.13 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 mandatory.

The following values are valid examples:

Q: loop

Q: process

Q: loop, discard

8.2.2.14 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.

8.2.2.15 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.

8.2.2.16 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

8.2.2.17 RestartMethod

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

RM: restart

8.2.2.18 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

8.2.2.19 MaxMGCPDatagram

The MaxMGCPDatagram parameter is encoded as a string of up to nine decimal digits – leading zeroes are not permitted. The following example illustrate the use of this parameter:

MD: 8100

8.3 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 18 below summarizes the response parameters whose presence is mandatory or optional in a response header, as a function of the command that triggered the response. 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 18 - Association of Response Header Parameters and Commands

Parameter name	CRCX	MDCX	DLCX	RQNT	NTFY	AUEP	AUCX	RSIP
ResponseAck ²⁵	O1	O1	O1	O1	O1	O1	O1	O1
CallId	F	F	F	F	F	F	O	F
ConnectionId	O2	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	O3	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
MaxMGCPDatagram	F	F	F	F	F	O	F	F
LocalConnection Descriptor	O4	O4	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.
2. In the case of a CreateConnection message, the response line is followed by a Connection-Id parameter and a LocalConnectionDescriptor. It may also be followed by a Specific-Endpoint-Id parameter, if the creation request was sent to a wildcarded Endpoint-Id. The connection-Id and LocalConnectionDescriptor parameter are marked as optional in the Table. In fact, they are mandatory with all positive responses, when a connection was created, and forbidden when the response is negative, and no connection was created.
3. Connection-Parameters are only valid in a successful response to a non-wildcarded DeleteConnection command sent by the Call Agent
4. A LocalConnectionDescriptor MUST be transmitted with a positive response (code 200) to a CreateConnection. It MUST also be transmitted in response to a ModifyConnection command, if the modification resulted in a change of the Local Connection Descriptor. The LocalConnectionDescriptor is encoded as a "session description", as defined in Section 8.4. It is separated from the response header by an empty line.

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

8.3.1 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 Section 8.4. 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 97 0
a=rtpmap:96 G726-32/8000
a=rtpmap:97 telephone-event/8000
a=mptime: 10 - 10
```

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 97 0
a=rtpmap:96 G726-32/8000
a=rtpmap:97 telephone-event/8000
a=mptime: 10 - 10
```

The final response is acknowledged by a Response Acknowledgement:

```
000 1204
```

8.3.2 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 Section 8.4. 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
a=mptime: 20
```

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
```

8.3.3 DeleteConnection

Depending on the variant of the DeleteConnection message, the response line may be followed by a Connection Parameters parameter line, as defined in Section 8.2.2.5.

```
250 1210 OK
P: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48,
PC/RPS=782, PC/ROS=45238, PC/RPL=5, PC/RJI=26
```

8.3.4 NotificationRequest

A NotificationRequest response does not include any additional response parameters.

8.3.5 Notify

A Notify response does not include any additional response parameters.

8.3.6 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 formatting restraints of this document.

```
200 1200 OK
A: a:PCMU; p:10, e:on, s:off, t:1, v:L,
m:sendonly;rcvonly;sendrecv;inactive
A: a:G728; p:20, e:on, s:off, t:1, v:L,
m:sendonly;rcvonly;sendrecv;inactive
A: a:G729; p:30, e:on, s:on, t:1, v:L,
m:sendonly;rcvonly;sendrecv;inactive;confrnce
```

8.3.7 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: CA-1@myhost.whatever.net:2345
```

```

L: mp:20;10, a:PCMU;G728
M: sendrecv
P: PS=622, OS=31172, PR=390, OR=22561, PL=5, JI=29, LA=50,
    PC/RPS=391, PC/ROS=22619, PC/RPL=5, PC/RJI=26

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 G728/8000
a=mptime: 10

```

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.

8.3.8 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@ca.whatever.net

```

8.4 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 [4].

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 "image", the session description is for an image service such as T.38 FAX relay.

8.4.1 Use of SDP

In an analog line gateway, we only have to describe sessions that use exactly one MIME media type at a given point in time; either "audio" (for voice or voice-band data) or "image" (for T.38 fax calls). The parameters of SDP that are relevant for both the "audio" and "image" based media types are specified in Section 8.4.2. Parameters specific to "audio" are specified in the Section 8.4.3. Parameters specific to "image" when used for T.38 are specific in section 8.4.4. Embedded clients MUST support session descriptions that conform to these rules and in the following order:

1. The SDP profile presented below
2. SDP: Session Description Protocol [4]

The CMS should be careful if it decides it is necessary to alter the SDP received from an endpoint. SDP provides a means of communicating the capabilities of an endpoint to another endpoint. If the CMS chooses to modify the SDP, it MUST NOT alter the SDP such that it violates the rules defined in this section.

The SDP profile provided describes the use of the session description protocol in NCS. The general description and explanation of the common individual parameters as well as most of the parameters being used for audio only can be found in [4]. Parameters specific to T.38 for image can be found in [27]. 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). It should be noted, that the SDP profile used here does not comply with the offer/answer model defined in RFC 3264. Thus, if a CMS needs to interact with another entity that uses the offer/answer model, the CMS may need to edit the SDP it receives from the endpoint.

8.4.2 SDP Parameters Common to both Audio (m=audio media line) and Image (m=image media line) Service Use

8.4.2.1 Protocol Version (v=)

v= <version>

"v= 0"

Send: MUST be provided in accordance with [4] (i.e., v=0).

Receive: MUST be provided in accordance with [4].

8.4.2.2 Origin (o=)

The origin field consists (o=) of 6 sub-fields in [4]:

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. ²⁷

Receive: This field SHOULD be ignored.

Session-ID:

Send: MUST be in accordance with [4] for interoperability with non-PacketCable clients.

Receive: This field SHOULD be ignored.

Version:

Send: Must be in accordance with [4].

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 [4] for interoperability with non-PacketCable clients.

Receive: This field MUST be ignored.

8.4.2.3 Session Name (s=)

s= <session-name>

"s= -"

Send: Hyphen MUST be used as Session name.

Receive: This field MUST be ignored.

8.4.2.4 Session and Media Information (i=)

i= <session-description>

²⁷ Since NCS endpoints do not know when privacy is requested, they should always use a hyphen.

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

Receive: This field MUST be ignored.

8.4.2.5 URI (u=)

u= <URI>

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

Receive: This field MUST be ignored.

8.4.2.6 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.

8.4.2.7 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.

8.4.2.8 Bandwidth (b=)

b= <modifier> : <bandwidth-value>

"b= AS : 64"

Send: Bandwidth information is optional in SDP, but it SHOULD always be included.²⁸ When an rtpmap or a non well-known codec is used, the bandwidth information MUST be used.²⁹

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

Modifier:

Send: Type 'AS' MUST be used.

²⁸ If this field is not used, the Gate Controller might not authorize the appropriate bandwidth.

²⁹ A non well-known codec is a codec not defined in the PacketCable codec specification [21].

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. When RFC 2198 [29] Redundancy is included in SDP, the bandwidth value included MUST include the maximum bandwidth required for use of redundancy with the specified depth.

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

Refer to the PacketCable Codec specification [11], for the details of calculating the bandwidth value.

8.4.2.9 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 [4].

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

8.4.2.10 Attributes (a=)

a= <attribute> : <value>

a= mptime: <alternative 1> <alternative 2> ...

a= rcvonly

a= sendrecv

a= sendonly

a= ptime

a=pmft: <modem-fax-transport>

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.

Note that SDP [4] requires unknown attributes to be ignored.

mptime:

This attribute is a media-level attribute defined by PacketCable. The mptime attribute defines a list of packetization period values the endpoint is capable of using (sending and receiving) for this connection.

Send: The mptime attribute MUST be present. There MUST be precisely one entry in the list for each <format> entry provided in the "m=" line. Entry number j in this list defines the packetization period for entry number j in the "m=" line. The first entry in the list MUST be a decimal number whereas subsequent entries in the list MUST be either a decimal number or a hyphen. For those media formats where a single packetization rate does not apply (e.g., non-voice codecs such as telephone-event or comfort noise), a hyphen ("-") MUST be encoded at the corresponding location in the list of packetization periods.

Receive: Conveys the list of packetization periods that the remote endpoint is capable of using for this connection; one for each media format in the "m=" line. For media formats whose packetization period is specified as a hyphen ("-"), the endpoint MUST use one of the packetization periods that was actually specified in the list. If the "mptime" attribute is absent, then the value of the "ptime" attribute, if present, MUST be taken as indicating the packetization period for all codecs present in the "m=" line.

recvonly:

Send: This attribute should not be supplied by an embedded client.

Receive: The field MUST be ignored.

sendrecv:

Send: This attribute should not be supplied by an embedded client.

Receive: The field MUST be ignored.

sendonly:

Send: This attribute should not be supplied by an embedded client.

Receive: The field MUST be ignored.

ptime:

Send: The ptime SHOULD be sent if it was received in a Remote Connection Descriptor or if the CMS used the packetization period ('p:') LocalConnectionOption.

Receive: The field MUST be ignored if the SDP contains the "mptime" attribute (as required in PacketCable compliant devices). If the "mptime" attribute is not present, then this field is used to define the packetization interval for all codecs present in the SDP description and the MTA MUST use the ptime in the calculation of QoS reservations.

pmft:

This is a session level attribute defined in ITU V.152 [30].

Send: The sending of the a=pmft line is NOT PERMITTED unless both the FXR and VBD packages are specified in the LCO by the CMS. An a=pmft: T38 line MUST be sent if V.152 is present in the Local Connection Descriptor and the LCO has one of the following two fxr/fx formats:

fxr/fx:t38-loose

fxr/fx:t38;gw.

If sent, the a=pmft line MUST be placed at the session level.

If any other fxr/fx format allowed in Section 7.7.5 is used, the endpoint MUST NOT send an a=pmft line.

An a=pmft line with modem-fax-transport other than T38 MUST NOT be sent.

Receive: The field MUST be ignored if present in the remoteconnectiondescriptor.

8.4.3 SDP Audio Service Use

The following SDP parameters are applied at the media level and are specific to Audio Service Use. Audio Service Use includes VBD mode as defined in V-152 [30]. PacketCable endpoints MUST NOT send any of these parameters within an image media descriptor (see Section 8.4.4). If, however, the endpoint receives an SDP with attribute parameters specific to image only within an audio media descriptor, the parameters SHOULD be ignored. Further, when media capability parameters are to be provided, each media capability descriptor (which includes the capability description line, a=cdesc, along with 0 or more capability attribute lines, e.g., a=cpar) MUST appear after the last media attribute and each media capability descriptor MUST be listed separately.

8.4.3.1 Encryption Keys:

k= <method>

k= <method> : <encryption-keys>

Security services for PacketCable are defined by the PacketCable Security specification [13]. The security services specified for RTP and RTCP do not comply with those of [2], [3], and [4]. In the interest of interoperability with non-PacketCable devices, the "k" parameter will therefore not be used to convey security parameters.

Send: MUST NOT be used.

Receive: This field SHOULD be ignored.

8.4.3.2 Attributes (a=)

a= <attribute> : <value>

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

a= rtpmap : 0 PCMU / 8000

a= fmp: <format> <format specific parameters>

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

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

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

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

a= X-pc-nrekey: <value>=<attribute>

a=gpm: <format> <parameter list>

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.

Note that SDP [4] requires unknown attributes to be ignored.

rtpmap:

Send: When used, the field MUST be used in accordance with [4]. This field MAY be used for well known as well as non well-known codecs. The encoding names used are provided in a separate PacketCable specification (see [11]). The mapping of codec to RTP dynamic payload type given with this attribute defines the payload type that this sender is prepared to receive on the connection. It also provides a strong hint to the other party that it should also use this payload mapping for its receive side, although there may be cases where this is not possible. On a given connection, once an MTA has mapped a dynamic payload type to a given encoding method for its receive media stream, that payload type MUST NOT subsequently be mapped to another encoding method for its receive media stream.

Receive: When used, the field MUST be used in accordance with [4]. This attribute defines the mapping of codec to RTP payload type that the other side of the connection is prepared to receive. MTAs MUST, therefore, use this payload type mapping when transmitting media on this connection. When received in a CreateConnection command, the MTA SHOULD use this payload type mapping for its own receive side (i.e., it should return a Local Connection Descriptor containing the same rtpmap attribute).

If an MTA receives a rtpmap attribute in a ModifyConnection command with a different mapping, the MTA MUST leave its own receive payload type mapping as is (so that asymmetric payload types are used).

fmp:

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 [19] 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 [4].

X-pc-codecs:

This attribute is a media-level attribute defined by PacketCable.

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 signaled through a media (m=) line.

X-pc-secret:

This attribute is a media-level attribute defined by PacketCable.

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 [4] with the following constraints:

- The encryption key 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 [17], 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 [13]. 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 [13] and it MUST be separated by at least one white space from the secret.

Pad and secret MUST use the same encode method.

X-pc-csuides-rtp:

X-pc-csuides-rtcp:

These attributes are media-level attributes defined by PacketCable.

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 [18]. Whitespaces are not allowed within a ciphersuite. The following example illustrates the use of ciphersuite: 62/51

The actual list of ciphersuites is provided in the PacketCable Security Specification [13].

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 signaled through a new ciphersuite line with the desired ciphersuite listed first.

X-pc-nrekey:

This attribute is a media-level attribute defined by PacketCable.

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 [13].

Receive: Conveys the number of rekey events. The field may be present when RTP security is used and its use is as defined in [13].

gpmd:

Send: The 'gpmd' (general-purpose media descriptor) attribute shall be used to associate payload types in a media information ('m') line with VBD mode.
a=gpmd:<format> <parameter list>

This field **MUST** only be sent with V.152 VBD declaration and in the context of VBD declaration, the <format> must be an RTP/AVP payload type. The <parameter list> **MUST** be a single "parameter=value" pair. For RTP/AVP formats, this parameter=value pair addresses a parameter that is not part of their standard MIME definition for the format specified. For sessions supporting V.152, the parameter **MUST** be the Boolean 'vbd' with the allowed values of 'yes' or 'no'. When set to 'yes' the attribute indicates that the implementation supports VBD mode as described in the vbd package.

Receive: The field **MUST** be ignored if it contains a parameter list other than "vbd=yes" or "vbd=no".
Absence of the 'gpmd' attribute with a "vbd=yes" attribute/value pair for any codec in the SDP session description **MUST** be construed as non support of VBD mode as described in the vbd package.

8.4.3.3 Media Announcements (m=)

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

m= <media> <port> <transport> <fmt list>
"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 [4]. The port specified is the receive port, regardless of whether the stream is unidirectional or bi-directional. The sending port may be different.

Receive: **MUST** be used in accordance with [4]. 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 [4] **MUST** be used. Specifically, this field contains a list of one or more RTP payload types that this MTA is prepared to receive on the connection and that it would prefer to send with. Each payload type is mapped uniquely to a codec, either statically or dynamically. The static mapping **SHOULD** be used if available (e.g., 0 for PCMU, 8 for PCMA). If a dynamic payload mapping is used, an RTPMAP attribute **MUST** also be present and the guidelines in Section 8.4.3.2 **MUST** be followed.

Receive: In accordance with [4]. Specifically, this indicates the payload type(s) that the other side of this connection is prepared to receive.

8.4.4 SDP Image Service Use for T.38

The following SDP parameters are applied at the media level and are specific to Image Service Use for T.38. PacketCable endpoints **MUST NOT** send any of these parameters within an audio media descriptor (see Section 8.4.3). If, however, the endpoint receives an SDP with attribute parameters specific to audio only within an image media descriptor, the parameters **SHOULD** be ignored. Further, when media capability parameters are to be provided, each media capability descriptor (which includes the capability description line, a=cpsc, along with 0 or more capability attribute lines, e.g., a=cpar) **MUST** appear after the last media attribute and each media capability descriptor **MUST** be listed separately.

8.4.4.1 Encryption Keys:

k= <method>

k= <method> : <encryption-keys>

There are currently no security services defined for the "image/t38" media type.

Send: MUST NOT be used.

Receive: This field **SHOULD** be ignored.

8.4.4.2 Attributes (a=)

a= <attribute> : <value>

a=T38FaxVersion: <version>

a=T38MaxBitrate: <bitrate>

a=T38FaxRateManagement: <faxratemanagement>

a=T38FaxMaxBuffer: <maxbuffer>

a=T38FaxMaxDatagram: <maxsize>

a=T38FaxUdpEC: <ECmethod>

a=T38FaxFillBitRemoval

a=T38FaxTranscodingMMR

a=T38FaxTranscodingJBIG

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. Attribute values are case-insensitive. Implementations **MUST** accept the lowercase, uppercase, and mixed upper/lowercase encodings of all attributes.

Note that SDP [4] requires unknown attributes to be ignored.

T38FaxVersion:

As defined in [11], support of T38 fax relay is limited to Group-3 rates corresponding with T38 versions 0,1 or 2.

As defined in [27]: The recipient of the offer **MUST** accept that version or modify the version attribute to be an equal or lower version when transmitting an answer to the initial offer. The recipient of an offer **MUST NOT** respond with an answer containing a higher version than that which was offered.

Also as defined in [27]: Early implementations of T.38 equipment may not provide a T.38 version number. In receipt of SDP without the version attribute, the endpoint **MUST** assume that the version is 0. This is applied in the following discussion on sending and receiving this attribute:

Send: The endpoint **MUST** indicate the version that it intends to use with the T38FaxVersion attribute. The endpoint **MUST** indicate the version which **MUST** be 0, 1 or 2. However, it **MUST NOT** indicate a version that is higher than the version received in a RemoteConnectionDescriptor.

Receive: If a RemoteConnectionDescriptor is received and the T38FaxVersion attribute is not included, then the endpoint MUST use version 0 of the T.38 specification. If the attribute is included, the endpoint MUST use a version of the specification that is the same or lower than the version indicated.

T38MaxBitRate:

Send: The T38MaxBitRate attribute MAY be included. If included, T38MaxBitRate MUST be set to 14400.

Receive: The T38MaxBitRate attribute SHOULD be ignored.

T38FaxRateManagement:

Send: The T38FaxRateManagement attribute MUST be included and MUST have a value of "transferredTCF" when UDPTL is used. With the value "transferredTCF", TCF is passed end-to-end as opposed to an attribute value of "localTCF" where TCF is generated locally. Note that "localTCF" is only appropriate when a reliable transport such as TCP is used.

Receive: When UDPTL is used, the T38FaxRateManagement attribute either MUST be present with a value of "transferredTCF" or it MUST be absent, in which case transferred TCF is assumed. All other values of the attribute MUST be rejected (error code 505 – Unsupported RemoteConnectionDescriptor)

T38FaxMaxBuffer:

Send: The T38FaxMaxBuffer attribute MUST NOT be included.

Receive: The T38FaxMaxBuffer attribute SHOULD be ignored.

T38FaxMaxDatagram:

Send: The T38FaxMaxDatagram attribute MUST be included. The value indicated MUST NOT be less than 160 bytes. This is based on 40ms packetization period and a 14,400 bps data rate. It includes the UDPTL datagram without the IP and UDP headers.

Receive: Endpoints MUST NOT send a datagram larger than that specified in the T38FaxMaxDatagram attribute. Prior to sending any T.38 datagram, the endpoint MUST ensure that is within the limits defined by this attribute. If the specified T38FaxMaxDatagram value is too small to support redundancy for a given datagram, but sufficient to support T.38 without redundancy, then the endpoint MUST send that T.38 datagram without redundancy. If the value is too small to allow the datagram to be sent without redundancy, the endpoint MUST NOT send the T.38 datagram and the endpoint MUST then generate a failure indication.

T38FaxUdpEC:

Support for redundancy is mandatory whereas support for forward error correction is optional. Use of either scheme requires negotiation.

Send: The T38FaxUdpEC attribute MUST be included. The value "t38UDPFEC" MAY be sent if FEC is supported and there either was no RCD provided with the command or the value of the attribute received in the RCD for this command is "t38UDPFEC". Otherwise "t38UDPRedundancy" MUST be sent.

Receive: Redundancy MUST be used if the value of the T38FaxUdpEC attribute is "t38UDPRedundancy". If the T38FaxUdpEC attribute is "t38UDPFEC" and FEC is supported by the endpoint, then FEC SHOULD be used. If the T38FaxUdpEC attribute is "t38UDPFEC" and FEC is not supported, then redundancy MUST be used. If this attribute is not included, the endpoint MUST NOT use redundancy or FEC.

T38FaxFillBitRemoval:

Support for fill bit removal is optional and any use of it needs to be negotiated.

Send: If fill bit insertion and removal is supported and desired to be used, and the command either did not include an RCD, or it included an RCD with the T38FaxFillBitRemoval attribute present, then T38FaxFillBitRemoval MUST be included and fill bit insertion and removal MUST then be used. In all other cases, the T38FaxFillBitRemoval attribute MUST NOT be included and fill bit insertion and removal MUST NOT be used.

Receive: Fill bit insertion and removal MUST NOT be used if the T38FaxFillBitRemoval attribute is absent.

T38FaxTranscodingMMR:

MMR transcoding does not apply to UDPTL-based T.38.

Send: When UDPTL is being used for T.38, the T38FaxTranscodingMMR attribute MUST NOT be included.

Receive: If the T38FaxTranscodingMMR attribute is present for UDPTL-based T.38, the command MUST be rejected (error code 505 – unsupported RemoteConnectionDescriptor).

T38FaxTranscodingJBIG:

JBIG transcoding does not apply to UDPTL-based T.38.

Send: When UDPTL is being used for T.38, the T38FaxTranscodingJBIG attribute MUST NOT be included.

Receive: If the T38FaxTranscodingJBIG attribute is present for UDPTL-based T.38, the command MUST be rejected (error code 505 – unsupported RemoteConnectionDescriptor).

8.4.4.3 Media Announcements (m=)

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

m= <media> <port> <transport> <fmt list>
 "m= image 3456 udptl t38"

Media:

Send: The 'image' media type MUST be used for UDPTL-based T.38.

Receive: The type received MUST be 'image' for UDPTL-based T.38.

Port:

Send: MUST be filled in accordance with [4]. The port specified is the receive port. The sending port may be different.

Receive: MUST be used in accordance with [4]. The port specified is the receive port. The sending port may be different.

Transport:

Send: The transport protocol 'udptl' MUST be used for UDPTL-based T.38.

Receive: The transport protocol MUST be 'udptl' for UDPTL-based T.38. Implementations SHOULD also tolerate the upper-case form "UDPTL", as well as mixed upper and lower-case forms of the string "udptl".

Media Formats:

Send: The media format MUST be "t38".

Receive: The media format MUST be "t38".

8.5 Transmission Over UDP

8.5.1 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 endpoint,³⁰ the commands **MUST** be sent to the default MGCP port, which is 2427 for commands sent to Gateways and 2727 for commands sent to Call Agent. To minimize backward compatibility issues, it is **RECOMMENDED** that the Call Agent always explicitly state the MGCP port to use in NCS messages (and not rely on the default). 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: If the command is a CreateConnection or ModifyConnection command a provisional response is sent, otherwise the command 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 Section 8.5.2.

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 Section 7.4.

8.5.2 Retransmission Strategy

This specification 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 (see, e.g., [8]):

- 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 seconds.

³⁰ Each endpoint may be provisioned with a separate Call Agent address and port.

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 milliseconds by default and the maximum value for the retransmission timer is 4 seconds by default. These default values may be altered by the provisioning process.

8.5.3 Maximum Datagram Size, Fragmentation and Reassembly

MGCP messages being transmitted over UDP rely on IP for fragmentation and reassembly of large datagrams. The maximum theoretical size of an IP datagram is 65535 bytes. With a 20-byte IP header, and an 8-byte header, this leaves us with a maximum theoretical MGCP message size of 65507 bytes when using UDP.

However, IP does not require a host to receive IP datagrams larger than 576 bytes [15] which would provide an unacceptably small MGCP message size. Consequently, MGCP mandates that implementations **MUST** support MGCP datagrams up to at least 4000 bytes, which requires the corresponding IP fragmentation and reassembly to be supported. Note, that the 4000 byte limit applies to the MGCP level. Lower layer overhead will require support for IP datagrams that are larger than this: UDP and IP overhead will be at least 28 bytes, and e.g., IPsec will add more as well.

It should be noted, that the above applies to both Call Agents and endpoints. Call Agents can audit endpoints to determine if they support larger MGCP datagrams than specified above. Endpoints do currently not have a similar capability to determine if a Call Agent supports larger MGCP datagram sizes.

8.6 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 piggybacked **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 Section 7.4.3.1).

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.

8.7 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 message³¹ 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.

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:

1. If $term_{new}$ does not contain any wildcards:
 - a. Unconfirmed responses to old commands where $term_{old}$ equals $term_{new}$.
 - b. Optionally, one or more unconfirmed responses where $term_{old}$ contained the "any-of" wildcard, and the endpoint-name returned in the response was $term_{new}$.
 - c. Optionally, one or more unconfirmed responses where $term_{old}$ contained the "all" wildcard, and $term_{new}$ is covered by the wildcard in $term_{old}$.
 - d. 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 $term_{old}$.
2. If $term_{new}$ contains the "all" wildcard:
 - a. Optionally, one or more unconfirmed responses where $term_{old}$ contained the "all" wildcard, and $term_{new}$ is covered by the wildcard in $term_{old}$.
3. If $term_{new}$ contains the "any of" wildcard:
 - a. 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.

³¹ As opposed to a command with a Response Acknowledgement parameter.

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.

8.8 Provisional Responses

In some cases, transaction completion times may be significantly longer than otherwise.³² 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.

In order to overcome this problem, a provisional response MUST therefore be issued if the transaction completion time is expected to exceed a small period of time (200ms is RECOMMENDED). 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. (Error Code 407 SHOULD be used.)

³² For instance, when resources are reserved and committed externally as part of a transaction.

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 ($T_{t_{longtran}}$). The purpose of this timer is primarily to detect endpoint failure. The default value of $T_{t_{longtran}}$ is 5 seconds, 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. Thus, the CMS **MUST** issue a "Response Acknowledgement" response whenever it receives a final response containing an empty "ResponseAck" parameter regardless of the receipt of a provisional response to the transaction since the provisional response may have been lost. 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 Section 7.4), 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.

9 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 provided in the PacketCable Security Specification [13] which should be consulted for further information.

Appendix A. Event Packages

This section defines an initial set of event packages for the various types of endpoints currently defined by PacketCable for embedded clients.

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	<p>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> <ul style="list-style-type: none"> • "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	<p>If nothing appears in this column for an event, then the event cannot be signaled on command by the Media Gateway Controller. Otherwise, the following symbols identify the type of event:</p> <ul style="list-style-type: none"> • "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.

A.1 Analog Access Lines

The following packages are currently defined for Analog Access Line endpoints:

- Line
- FAX
- VoIP Metrics Package
- Media Format Parameter Package
- Voice Band Data Package
- General Purpose Media Descriptor Package

A.2 Line Package

Package name: L

The codes listed in Table 19 below are used to identify events and signals for the "line" package for "analog access lines":

Table 19 - Line Package Codes for Events and Signals

Code	Description	Event	Signal	Additional Info
0-9,*,#,A, B,C,D	DTMF tones	√	BR	
bz	Busy tone	-	TO	Time-out = 30 seconds.
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 seconds.
ft	Fax tone	√	-	
hd	Off-hook transition	P, S	-	
hf	Flash hook	P	-	
hu	On-hook transition	P, S	-	
L	DTMF long duration	√	-	
ld	Long duration connection	C	-	
ma	Media start	C	-	
mt	Modem tones	√	-	
mwi	Message waiting indicator	-	TO	Time-out = 16 seconds.
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 seconds.
rg	Ringing	-	TO	Time-out = 180 seconds.
ro	Reorder tone	-	TO	Time-out = 30 seconds.
rs	Ringsplash	-	BR	
rt	Ring back tone	-	C, TO	Time-out = 180 seconds.
sl	Stutter dial tone	-	TO	Time-out = 16 seconds.
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 = [(MaxReps + 1) + (MaxReps * Delay)] seconds, where the default values for MaxReps = 1 and Delay = 10 as defined in the NCS Signalling MIB specification.
X	DTMF tones wildcard	√	-	Matches any of the digits "0-9".

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

DTMF tones (0-9,*,#,A,B,C,D): Detection and generation of DTMF signals is described in GR-506-CORE – LSSGR: SIGNALING [20], Section 15. It is considered an error to try and play DTMF tones on a phone that is 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 a combination of two AC tones with frequencies of 480 and 620 Hertz and levels of –24 dBm each, to give a combined level of –21 dBm. The cadence for Station Busy Tone is 0.5 seconds on followed by 0.5 seconds off, repeating. See [20], Section 17.2.6. It is considered an error to try and play busy tone on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Confirmation tone (cf): Confirmation Tone uses the same frequencies and levels as dial tone (350 and 440 Hertz) but with a cadence of 0.1 second on, 0.1 second off repeated three times. See [20], Section 17.2.4. It is considered an error to try and play confirmation tone on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Caller Id (ci(time, number, name)): See TR-NWT-001188, [20], and TR-NWT-000031. Each of the three fields are optional, however each of the commas will always be included.

- The **time** parameter is coded as "MM/DD/HH/MM", where the first 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(10/14/17/26, "555 1212", CableLabs)
```

In addition to the generic signaling requirements described in Section 8.2.2.10, the MTA MUST support, at a minimum, the combination of one signal on an endpoint along with CallerId (ci) within the same SignalRequest line (e.g., S: rg, ci(time,number,name)) and a signal on each connection associated with the endpoint.

Dial-tone (dl): Dial Tone is a combination of two continuous AC tones with frequencies of 350 and 440 Hertz and levels of –13dBm each to give a combined level of –10 dBm. See [20], Section 17.2.1. It is considered an error to try and play dial-tone on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Fax tone (ft): The fax tone event is generated whenever a fax call is detected by presence of V.21 fax preamble. or T.30 CNG tone (if provisioned). See PacketCable 1.5 Audio/Video Codecs Specification [11], and ITU-T Recommendations T.30 [21] and V.21 [23].

Off-hook transition (hd): See [20], Section 12.

Flash hook (hf): See [20], Section 12.

On-hook transition (hu): See [20]. The timing for the on-hook signal is for flash response enabled.

DTMF Long duration (L): The "DTMF Long duration" is observed when a 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 DTMF signal, and then, 2 seconds 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 is 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 valid³³ 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.

The event is 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 [24] and V.8 [22].

Message Waiting Indicator (mwi): Message Waiting indicator tone uses the same frequencies and levels as dial tone (350 and 440 Hertz at -13dBm each) but with a cadence of 0.1 second on, 0.1 second off repeated 10 times followed by steady application of dial tone. See [20], Section 17.2.3. It is considered an error to try and play message waiting indicator on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Open Switch Interval (osi): See [20]: Voiceband Data Transmission Interface, Section 2.2.2. See also [20], Section 4.5.2.1.

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 dialed 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/dl)

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:

O: L/oc(B/C)

Note the use of "B" above 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.

³³ When authentication and integrity security services are used, an RTP packet is not considered valid until it has passed the security checks.

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

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

Note the use of "B" above as the prefix for the parameter reported.

Off-hook warning tone (ot): Receiver Off Hook Tone (ROH Tone) is the irritating noise a telephone makes when it is not hung up correctly. ROH Tone is generated by combining four tones at frequencies of 1400 Hertz, 2060 Hertz, 2450 Hertz and 2600 Hertz at a cadence of 0.1 second on, 0.1 second off, repeating. GR-506-CORE [20], Section 17.2.8 contains details about required power levels. It is considered an error to try and play off-hook warning tone on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Distinctive ringing (r0, r1, r2, r3, r4, r5, r6 or r7): See [20], Section 14. The values for r1 to r5 are as defined for distinctive ringing pattern 1 to 5 as defined in [20]. The provisioning process MAY define the ringing cadence for each of the r0-r7 signals. The MTA MUST support provisioning for r0, r6 and r7. It is considered an error to try and ring a phone that is off hook and an error should consequently be returned when such attempts are made (error code 401 – phone off hook).

Ringling (rg): See [20], Section 14. The value for rg is as defined for distinctive ringing pattern 1 as defined in [20]. The provisioning process MAY define the ringing cadence. The ringing signal may be parameterized with the signal parameter "rep" which specifies the maximum number of ringing cycles (repetitions) to apply. The following will apply the ringing signal for up to 6 ringing cycles:

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

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

Reorder tone (ro): Reorder tone is a combination of two AC tones with frequencies of 480 and 620 Hertz and levels of –24 dBm each, to give a combined level of –21 dBm. The cadence for reorder tone is 0.25 seconds on followed by 0.25 seconds off, repeating continuously. See [20], Section 17.2.7. It is considered an error to try and play reorder tone on a phone that is on hook and an error should consequently be returned when such attempts are made (error code 402 – phone on hook).

Ringsplash (rs): Ringsplash, also known as "Reminder ring," is a burst of ringing 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. In the US, it is defined to be a 0.5(-0,+0.1) second burst of power ringing. See TR-TSY-000586 – Call Forwarding Subfeatures. The provisioning process MAY define the cadence for ringsplash. It is considered an error to try and ring a phone that is off hook and an error should consequently be returned when such attempts are made (error code 401 – phone off hook).

Ring back tone (rt): Audible Ring Tone is a combination of two AC tones with frequencies of 440 and 480 Hertz and levels of –19 dBm each, to give a combined level of –16 dBm. In the U.S. the cadence for Audible Ring Tone is defined to be 2 seconds on followed by 4 seconds off. The definition of the tone is defined by the national characteristics of the Ringback Tone, and MAY be established via provisioning. See [20], Section 17.2.5.

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 on hook and an error should consequently be returned when such attempts are made (error code 402 – phone 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 generated by supplying Confirmation Tone, followed by continuous Dial Tone. See [20], Section 17.2.2. The stutter dial tone signal may be parameterized with the signal parameter "del" which will specify a delay in milliseconds to apply between the confirmation tone and the dial tone³⁴. The following will apply stutter dial tone with a delay of 1.5 seconds between the confirmation tone and the dial tone:

³⁴ This feature is needed for, e.g., Speed Dialing.

S: sl(del=1500)

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

Timer (t): As described in Section 7.1.5, timer T is a provisionable timer that can only be cancelled by 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, 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 seconds and the default value for T_{crit} is 4 seconds. 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 [25].

Visual Message Waiting Indicator (vmwi): The transmission of the VMWI messages will conform to the requirements in TR-H-000030 – Section 2.3.2, "On-hook Data Transmission Not Associated with Ringing", and the CPE guidelines in SR-TSV-002476. 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. When the endpoint restarts, it should not attempt to turn off a visual message waiting indicator, even if it believes it should be off. The Call Agent should periodically refresh the CPE's visual indicator. See TR-NWT-001401 – Visual Message Waiting Indicator Generic Requirements; and [20] Voiceband Data Transmission Interface.

Call Waiting tone1 (wt1, ..., wt4): Call Waiting tones are defined in [20], Section 14.2 – the number refers to the tone pattern used. The duration of the tone pattern lasts for up to 1 second. GR-571-CORE (FSD 01-02-1201) indicates that two tone patterns should be played separated by a period of 10 seconds. The default maximum number of repetitions and delay between repetitions may be altered by the provisioning process. The default Call Waiting tone is a 440-Hz tone applied for 300 ± 50 ms. The talking path should be interrupted for a maximum of 400 ms for the application of each CW tone pattern. When this signal is requested, the embedded client will play the two tone patterns as specified in [20] before the "TO" signal times out. It is considered an error to try and apply call waiting tones on a phone that is on hook and an error SHOULD consequently be returned when such attempts are made (error code 402 – phone on hook).

The default time-out is calculated based on the default (or provisioned) number of repetitions, default (or provisioned) delay between repetitions and length of time to generate the call waiting tone itself (1 second). For example, if Call Waiting Delay is set to 12 seconds and Call Waiting Maximum Number of Repetitions is set to 2, then the call waiting tone will be heard up to three times (played once and repeated twice). Therefore, the default time-out is calculated by adding the length of time it takes to play the call waiting tone three times (3 seconds) and the delay between repetitions ($2 * 12$ seconds). The calculated default timeout in this example is therefore 27 seconds. Upon request the call waiting signal may be parameterized with the signal parameter "to" which in essence overrides the calculated default timeout described above.

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

A.3 FAX Package

The following Media Gateway Control Protocol Fax Package defines a mechanism to support FAX calls. NCS imposes the following additional requirements above and beyond those specified in the Media Gateway Control Protocol Fax Package:

- Implementations **MUST** conform to the NCS Protocol Version encoding as specified in Section 8.2.1.4.
- The endpoint **MUST** be capable of generating the "t38(start)" event when T.30 CNG tone is detected. Furthermore, the endpoint **MUST** provide a configuration option to disable this capability.

Media Gateway Control Protocol Fax Package

Abstract

This document defines a Media Gateway Control Protocol (MGCP) package to support fax calls. The package allows for fax calls to be supported in two different ways. The first one utilizes ITU-T Recommendation T.38 for fax relay under the control of the Call Agent. The second one lets the gateway decide upon a method for fax transmission as well as handle the details of the fax call without Call Agent involvement.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC-2119 [RFC2119].

Table of Contents

1. Introduction
2. Fax Package Definition
2.1 LocalConnectionOptions
2.1.1 T.38 Procedure (Strict or Loose)
2.1.2 Gateway Procedure
2.1.3 Off Procedure
2.1.4 Mode Operation
2.1.5 Detecting a Fax Call
2.1.6 Considerations for Which Procedures to Request
2.2 Events and Signals
2.2.1 Gateway Controlled Fax (gwfax)
2.2.2 No Special Fax Handling (nopfax)
2.2.3 T.38 fax relay(t38):
2.3 Connection Parameters
2.4 Negotiation of T.38 Parameters
2.5 Implementation Considerations
2.5.1 Media IP Address and Port for T.38
2.5.2 Case Sensitivity
2.5.3 Boolean Indicator after T.38 Parameters
3. Call Flow Examples
3.1 Call Agent Controlled T.38 Strict
3.2 Multiple and Different Options
3.3 Interaction with SIP Endpoints
4. Security Considerations

5. IANA Considerations
6. Normative References
7. Informative References
8. Acknowledgements
9. Author's Address

1. Introduction

This document defines a Media Gateway Control Protocol (MGCP) [RFC3435] package that enables MGCP controlled gateways to support fax calls. The package enables fax calls to be supported in two different ways. The first one utilizes ITU-T Recommendation T.38 using either UDPTL or TCP (see [T38]) for fax relay under the control of the Call Agent. The second one lets the gateway decide upon a method for fax transmission as well as handle the details of the fax call without Call Agent involvement.

The fax package definition is provided in Section 2 and in Section 3 we provide two call flow examples showing how to use it. Security considerations are found in Section 4, followed by the IANA considerations and references.

2. Fax Package Definition

A package is defined for fax. The package defines new localconnectionoptions, events, and connection parameters as detailed below.

Package Name: FXR

Package Version: 0

2.1 LocalConnectionOptions

A new Fax LocalConnectionOptions (LCO) parameter is defined for fax handling. The Call Agent supplies this fax LCO to indicate the desired fax handling procedure to the Media Gateway. The fax LCO contains a list of desired fax handling procedures ordered by preference with the most desired procedure listed first. When the parameter is explicitly included in a command, the gateway MUST be able to use at least one of the listed procedures for the command to succeed. The list can currently indicate one or more of the following procedures (see Section 2.1.1 to 2.1.4 for further details on these):

* T.38 Strict Use T.38 [T38] with either UDPTL or TCP for fax relay and have the Call Agent control it. Assuming the procedure can be used (see Section 2.1.1), a switch to T.38 procedures will be initiated upon fax detection and a "t38(start)" event will be generated (see Section 2.2). This mode requires an indication of T.38 support from the remote side in order to be used, as described further in Section 2.1.1.

* T.38 Loose Identical to T.38 Strict procedure, except that an indication of T.38 support from the remote side is not required for the procedure to be used.

* Off Do not invoke any special procedure for fax, except for echo cancellation adjustment and possibly switching to another codec.

* Gateway Let the gateway control and decide how to handle fax calls without Call Agent involvement. This includes the case where the gateway does not do

anything special for fax, hence by definition this procedure can always be supported. If the gateway does invoke a special procedure upon detection of fax, it will generate a "gwfax(start)" event so the Call Agent can be notified about it (see Section 2.2). The Call Agent SHOULD then refrain from issuing potentially conflicting commands to the gateway until the gateway ends its special fax handling procedure.

A gateway that ends up not being able to invoke any special procedure for fax will generate a "nopfax(start)" event (see Section 2.2) upon detection of fax.

The set of possible values (i.e., procedures) for the fax LCO is extensible. The prefix "x-", which indicates an optional extension, and the prefix "x+", which indicates a mandatory extension, are reserved for vendor-specific use.

In CreateConnection commands, the fax LCO value defaults to "gateway". In ModifyConnection commands, the fax LCO value defaults to its current value on the connection. Thus, if LocalConnectionOptions are either omitted or the fax LCO is not included in a ModifyConnection command, the previous fax LCO value for the connection is retained, but without affecting the outcome of the command; consequently, the gateway may now not apply any special procedure to fax. If the Call Agent wants to ensure that a command succeeds only when a fax procedure is applied, the command needs to include the fax LCO explicitly.

As an example of this, assume that the CreateConnection command successfully specified the use of "T.38 Strict", and a ModifyConnection command is now received without the fax LCO, but with a RemoteConnectionDescriptor indicating no support for T.38; in that case, the ModifyConnection command will succeed, however T.38 procedures will no longer be invoked upon fax detection. Had the Call Agent instead included the fax LCO set to "T.38 Strict", the command would have failed.

If multiple fax parameter values are provided, the gateway MUST choose one of the procedures specified according to the order in which they are supplied, except as follows:

1. If "gateway" would have been selected and it would have resulted in no special procedure being applied, and
2. there are other procedures than "off" which are specified after "gateway" (e.g., "t38")

then the gateway MUST use the most preferred of those subsequent procedures that can be supported. If none of those subsequent procedures can be supported, the gateway reverts to not invoking any special procedure for fax. Please refer to Section 2.1.4 for further details on determining which procedures can be supported.

The fax LCO parameter is encoded as the keyword "fx" (prefixed with the package name per [RFC3435]), followed by a colon and a semicolon separated list of values where T.38 Strict is encoded as "t38", T.38 Loose is encoded as "t38-loose", gateway is encoded as "gw", and off is encoded as "off".

The following example illustrates use of PCMU or G.729 for audio encoding and T.38 Strict fax relay (preferred) or gateway control for fax:

```
L: a:PCMU;G729, fxr/fx:t38;gw
```

It should be noted, that MGCP allows the CreateConnection command to omit both LocalConnectionOptions and a RemoteConnectionDescriptor thereby letting the gateway decide upon the media parameters to use. When the T.38 fax package is supported, the gateway thus could choose to do either audio or T.38 fax relay in such cases. Most likely, the Call Agent requires one or the other to be used, and hence it SHOULD NOT omit both LocalConnectionOptions and a RemoteConnectionDescriptor in CreateConnection commands.

When auditing capabilities, the fax LCO may be returned with a semi-colon separated list of supported fax handling parameters. The values "t38-loose", "off" and "gw" MAY be omitted from such a list as they are always implied. Gateways that implement additional parameters SHOULD return these additional parameters when capabilities are audited as illustrated by the following example:

```
A: a:image/t38, fxr/fx:mypar, ...
```

In the following subsections we provide additional detail on the above defined fax procedures.

2.1.1 T.38 Procedure (Strict or Loose)

When a gateway is instructed to use one of the T.38 procedures (strict or loose), also known as Call Agent controlled T.38 mode, the "m=" line in the SDP returned will not indicate use of UDPTL-based or TCP-based T.38 (unless the gateway was also instructed to use "image/t38" for the media stream). Any other entity seeing this SDP will not know whether T.38 is supported or not and hence whether it is safe to attempt a switch to T.38 upon fax detection. To remedy this dilemma, capability information for T.38 (if supported) using the SDP Simple Capability Declaration extensions [RFC3407] MUST be included. Other capability information is included as well, regardless of whether the Call Agent authorized use of those in the connection handling command or not. A subsequent attempt to actually use these may of course not succeed, e.g., because the Call Agent LCOs does not allow them to be used. The following example illustrates the RFC 3407 capability descriptor - note the inclusion of both current (audio) and latent (T.38) capabilities as specified in RFC 3407:

```
m=audio 3456 RTP/AVP 18
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 18
a=cdsc: 2 image udptl t38
```

For a list of T.38 related parameters to be included in the SDP, please refer to T.38 Annex D [T38].

Upon fax detection, a gateway which has successfully been instructed to use one of the T.38 procedures will:

1. Initiate the T.38 fax relay procedure and mute the media channel in both the send and receive direction (unless the media channel is already using T.38).
2. Generate a "t38(start)" event.
3. Await further instructions from the Call Agent in order to initiate the actual media change (unless the media channel is already using T.38).

The Call Agent instructs the gateway to perform the media change by sending it a ModifyConnection command with "image/t38" listed as the encoding method in the LocalConnectionOptions (receipt of a ModifyConnection command without LocalConnectionOptions but with a RemoteConnectionDescriptor containing an "m=" line with the MIME type "image/t38" would achieve the same). Per the normal MGCP codec negotiation procedures (see [RFC3435] Section 2.6), if a RemoteConnectionDescriptor was included as well, it needs to include an "m=" line with "image/t38" as an acceptable media format in order for the command to succeed. The gateway may choose between the UDPTL and TCP transport protocols at its own discretion subject to the normal MGCP codec negotiation procedures (in practice, TCP-based implementations are currently rare).

If a RemoteConnectionDescriptor was not included with the ModifyConnection command sent to a gateway that initiated the T.38 procedure, it is possible (in fact likely), that the last received RemoteConnectionDescriptor did not include an "m=" line listing "image/t38" as an acceptable media format. In that case, the endpoint cannot send T.38 media to the other side. The endpoint MUST instead wait for an updated RemoteConnectionDescriptor with "image/t38" as an acceptable media format and a supported transport protocol (UDPTL or TCP). The T.38 fax procedure continues when an acceptable RemoteConnectionDescriptor is received. An acceptable RemoteConnectionDescriptor contains an "m=" line with the "image/t38" MIME type (using the normal SDP syntax) and a supported transport protocol (UDPTL or TCP). If the fax call fails, e.g., due to a fax timeout, while either waiting for the Call Agent to instruct the gateway to switch to "image/t38" or waiting for an acceptable RemoteConnectionDescriptor, a "t38(stop)" or a "t38(failure)" event MUST be generated. When the T.38 procedure ends, a "t38(stop)" or "t38(failure)" event MUST be generated.

Finally, the Call Agent may need to abort a T.38 procedure that is in progress. This can for example be done when the remote side was unable to switch to T.38, and a fallback to fax passthrough using an audio codec is attempted. The Call Agent instructs the endpoint to abort an in-progress T.38 procedure by use of the "off" Fax LCO as illustrated below:

```
L: fxr/fx:off
```

We now define "time t38init" as the point in time where the T.38 procedure was initiated, and "time t38abort" as the point in time where the Call Agent aborts an in-progress T.38 procedure. If the Call Agent at time t38abort instructs or enables the endpoint to revert to one or more codecs that were in use just prior to time t38init, the endpoint SHOULD use media stream parameters that mimic the most recent LocalConnectionDescriptor issued before time t38init. For example, IP-address and UDP port, payload formats used and their payload type mapping, should all be the same as before time t38init. This will enable the fallback to be as rapid as possible. A LocalConnectionDescriptor is returned as usual, i.e., only if one or more parameters changed since the last LocalConnectionDescriptor issue (e.g., if a T.38 LCD was issued or a transport address in the audio LCD was changed).

2.1.2 Gateway Procedure

A gateway using the gateway procedure, also known as Gateway controlled mode, may initiate special fax handling upon detecting a fax call. The details of this special fax handling are outside the scope of this document. However, in order to use any special fax handling, support for it MUST be negotiated with the other side by passing

and recognizing relevant parameters via the LocalConnectionDescriptor and RemoteConnectionDescriptor. If the other side has not indicated support for the special fax handling desired, the gateway MUST NOT attempt to initiate it. When special fax handling is initiated, a "gwfax(start)" event MUST be generated thereby enabling the Call Agent to differ between the Call Agent and gateway controlled mode while still being informed about the actual change to fax. When the special gateway handling of fax ends, a "gwfax(stop)" or "gwfax(failure)" event MUST be generated.

2.1.3 Off Procedure

A gateway using the "off" procedure will not invoke any special fax procedures, e.g., T.38, when detecting a fax. However, the gateway may still adjust local echo cancellation and/or switch to an alternative codec as needed (in particular, this does not preclude the use of RTP-based T.38). Also, a "nopfax(start)" event MUST be generated; a corresponding "stop" event however will not.

Generating a "stop" event would imply that the gateway had to infer when the fax call ends, which involves processing of the media stream. However, when using the "off" mode, such processing is not expected to occur.

2.1.4 Mode Operation

For each of the above modes, the RemoteConnectionDescriptor provides information on what procedure(s) the other side supports. The following rules are used to determine which procedure to use:

1. Whatever the Call Agent specified in the Fax LocalConnectionOptions for the current command MUST be adhered to. If the gateway cannot satisfy any of the options, the command fails (error code 532 - unsupported value(s) in LocalConnectionOptions is RECOMMENDED).

2. If both Fax LocalConnectionOptions and a RemoteConnectionDescriptor are provided, the procedure selected MUST be supported by both sides - this is currently only an issue for "T.38 Strict". A procedure can be satisfied by the remote side if:

- * the relevant MIME media type, e.g., "image/t38", is included in the "m=" line in the RemoteConnectionDescriptor, or

- * the relevant MIME media type is included as a capability (see [RFC3407]) in the RemoteConnectionDescriptor.

If the gateway cannot select any of the procedures in the Fax LocalConnectionOptions, the command fails (error code 532 is RECOMMENDED). Note that "T.38 Loose", "gateway", and "off" by definition can always be supported by an implementation that supports this package, irrespective of what the RemoteConnectionDescriptor indicates.

3. If the Call Agent did not include any Fax LocalConnectionOptions or a RemoteConnectionDescriptor with the command, the gateway MUST continue using whichever procedure it is currently using.

4. If the Call Agent did not include any Fax LocalConnectionOptions, but a RemoteConnectionDescriptor was included, the gateway MUST follow rule 2 in selecting a procedure. In so doing, the default Fax LocalConnectionOptions, i.e., "gateway" in

CreateConnection, or the current value in ModifyConnection, MUST be used. In the case of ModifyConnection, the outcome of the command does not depend on the gateway being able to select one of these "default" procedures (as described in Section 2.1). Note that this is not an issue for the CreateConnection command, since the default value can always be supported by definition.

5. A previously received RemoteConnectionDescriptor does not affect what procedure can be selected. Only a RemoteConnectionDescriptor supplied with the current command affects the procedure selection. However, in order to send media of a given type (e.g., "image/t38"), the most recently received RemoteConnectionDescriptor MUST include a corresponding media line.

The following examples illustrate the use of the above rules:

Per rule 1, a gateway that only supports standard T.38 fax relay will fail a command that only contains the fax option "mypar" whereas it will succeed a command that contains, "t38-loose", "gw", "off" or no Fax LCO. A command that only contained "t38", i.e., use of T.38 in "strict" mode, may or may not succeed (depending on the RemoteConnectionDescriptor).

A gateway supporting T.38 that receives a CreateConnection command with the fax handling LCO set to "t38" and a RemoteConnectionDescriptor with neither a T.38 capability nor a T.38 media stream will fail per rule 2. Had the fax handling LCO included either "t38-loose", "gw" or "off", the command would have succeeded and any of the procedures included could have been selected.

Assume a gateway supporting T.38 has successfully executed a CreateConnection command with fax handling set to "t38" (i.e., strict). If the gateway now receives a ModifyConnection command without a fax handling LCO but with a RemoteConnectionDescriptor that has neither a T.38 capability nor a media stream with "image/t38", the command will succeed (since rule 1 has no effect in that case). However, per rule 2 and 4, there will not be any T.38 procedure in place. Had the CA instead included a fax handling LCO set to "t38" again, the command would have failed per rule 2.

Finally, it should be noted that a switch to T.38 can be initiated by either one or both of the originating and terminating gateways and hence implementations MUST be prepared to handle this. This includes the case where both sides initiate the switch, which for example can occur when the originating fax generates Calling Tone (CNG) and the terminating fax detects V.21 fax preamble (see [T30]) before the switch to T.38 has been performed on the terminating side.

2.1.5 Detecting a Fax Call

A fax call can be detected by several different means, e.g., V.21 fax preamble, T.30 CNG tone, or V.8 signals, depending on the fax transmission method being used. Implementations of this package MUST at a minimum detect a fax call based on V.21 fax preamble.

Triggering based on T.30 CNG tone MAY be done; this is generally considered acceptable for G3 and lower fax speeds. However, when used with T.38 version 2 or earlier, it will impact V.34 high-speed fax. The reason is, that T.38 version 2 (and earlier) does not support the V.8 ANSam and CM signals used with V.34 fax, and hence the V.34 faxes will downspeed to G3 (14.400 bps) or lower when using T.38 version 2 (or

earlier). Also, a few rare cases of modems generating T.30 CNG tones for non-fax call have been reported; such modems would generate a false trigger for fax. As a consequence of the above, it is RECOMMENDED that implementations of this package which support T.30 CNG based fax detection provide a configuration option to disable it for T.38 version 2 (or earlier).

2.1.6 Considerations for Which Procedures to Request

It is important to understand the implications of using any one of the above defined procedures. Furthermore, multiple alternative procedures can be requested, however not all combinations make sense. In this section, we elaborate on both of these issues.

Use of the T.38 strict mode is ideal in an environment where it is known, that other endpoints generate RFC 3407 capability descriptions with T.38 fax relay information. If a RemoteConnectionDescriptor without T.38 fax relay capabilities is received in such an environment, it is known that the other side does not support T.38, and hence an unsuccessful attempt to switch to T.38 (which in turn may lead to a failed fax call) can be avoided. If it is not known whether other endpoints support the RFC 3407 capability descriptors or not, the tradeoff is less clear. The advantage is, that a switch to T.38 will only be attempted if it is known that the other side supports it, however endpoints that do not indicate support for T.38 may still support it; T.38 however will not be used with these which in turn may lead to unnecessary fax failures with low-bandwidth codecs or lossy networks.

Use of the T.38 loose mode involves the same considerations as for T.38 strict, however the pros and cons are reversed. If a peer endpoint does not support T.38, the T.38 loose mode will still attempt to switch to T.38 (and fail), which in turn may lead to a failed fax call. On the other hand, if the peer endpoint does not support the RFC 3407 capability descriptors, but the peer endpoint does in fact support T.38, T.38 would still be used with this mode.

In summary, there is no single good answer to the use of either T.38 strict or T.38 loose mode; it depends on the capabilities of the endpoints involved as well as the trade-off between potentially letting fax calls fail due to lack of capability indications (where T.38 otherwise is supported) versus potentially letting fax calls fail due to an unsuccessful switch to T.38 (because T.38 in fact was not supported). It should be noted, that Call Agents may have means beyond RFC 3407 capability descriptors to determine if a peer endpoint supports T.38 or not. For example, when SIP is used as the signaling protocol with other peers (e.g., Call Agents or other SIP devices), the SIP OPTIONS method can be used to learn whether T.38 is supported. Also, if the Call Agent allows use of high-bandwidth codecs with redundancy when support for T.38 is not indicated, fax calls may still succeed without the use of T.38, even in networks with non-negligible packet loss.

When the gateway controlled mode is selected, there will only be special fax handling if the two peer endpoints support the same fax handling method; note that the details of the actual method is entirely up to the vendor. Also note, that if the two peer endpoints either do not support the same method for fax handling, or the method is not indicated in the SDP exchanged, there will be no special fax handling in place. Furthermore, the Call Agent will not be aware that this is the case until the fax transmission starts and a "nopfax(start)" event is generated.

The off mode is straightforward; there will be no special procedure in place for fax handling, except for the usual handling of echo cancellation and possibly a change to a higher bandwidth codec.

Having looked at the individual procedures in more detail, we now elaborate on some of the combinations of procedures that may be requested:

* T.38 strict: If the T.38 strict procedure is placed after the T.38 loose or the off procedure (both of which can always be supported), it will not be selected. Apart from this, it makes little sense to request both T.38 strict and T.38 loose.

* T.38 loose: The T.38 loose procedure can always be supported, so any procedure specified after T.38 loose will not be selected.

* Gateway: The gateway controlled procedure can always be supported. If the gateway controlled procedure would have resulted in no special fax procedure and further options (except off) are provided, those procedures will be attempted. If neither of those procedures can be supported, there will be no special fax procedure in place.

* Off: The off procedure can always be supported. Any procedure specified after this one will not be selected.

2.2 Events and Signals

The following events are defined in support of the above:

Symbol	Definition	R	S	Duration
gwfax	Gateway controlled fax	x		
nopfax	No special fax handling	x		
t38	T.38 fax relay	x		

The definitions of the individual events are provided in the following subsections.

2.2.1 Gateway Controlled Fax (gwfax)

The "gateway controlled fax" event occurs when the gateway handled fax procedure either starts, stops or fails. The event is encoded as "gwfax" and the following event parameters, which apply to ObservedEvents only, are defined:

* start Gateway controlled fax procedure was initiated. The Call Agent SHOULD refrain from issuing media handling instructions to the gateway until either a "gwfax(stop)" or "gwfax(failure)" event is generated.

* stop Gateway controlled fax procedure ended and the gateway did not detect any errors. Note that this does not necessarily imply a successfully transmitted fax. It merely indicates that the gateway controlled fax procedure has ended and the procedure itself did not encounter any errors. Media parameters for the connection are as before the gateway handled fax procedure started.

* failure The gateway controlled fax procedure ended abnormally. Some kind of problem was encountered in the gateway controlled fax procedure and the procedure ended. Media parameters are as before the gateway handled fax procedure started.

One of the above parameters will be present when the event is reported. The "gwfax" event MAY be parameterized with additional parameters in ObservedEvents, however it is RECOMMENDED that one of the above parameters is the first parameter supplied. Unknown parameters MUST be ignored.

The following example illustrates the encoding of the "gwfax" event:

```
O: fxr/gwfax(start)
O: fxr/gwfax(stop, foobar)
```

2.2.2 No Special Fax Handling (nopfax)

The "no special fax handling" event occurs when there is no special fax handling procedure in place and a fax call is detected. This can happen either due to no special fax handling procedure being requested (including "off"), or negotiation leading to no special fax handling procedure being possible. The event is encoded as "nopfax" and the following event parameter, which applies to ObservedEvents only, is defined:

* start No special fax handling procedure is in place, however a fax call is now detected. The Call Agent may have to issue further commands in order to ensure a successful fax call (e.g., switch to another codec).

The above parameter will be present when the event is reported. The "nopfax" event MAY be parameterized with additional parameters on ObservedEvents, however it is RECOMMENDED that the above parameter is the first parameter supplied. Unknown parameters MUST be ignored. Note, that this event currently cannot be parameterized with "stop" or "failure" as it only detects the beginning of a fax call.

The following example illustrates the encoding of the "nopfax" event:

```
O: fxr/nopfax(start)
```

2.2.3 T.38 fax relay(t38):

The "T.38 fax relay" event occurs when one of the T.38 fax relay procedures (strict or loose) either starts, stops or fails. The event is encoded as "t38" and the following event parameters, which apply to ObservedEvents only, are defined:

* start A fax call was detected on the endpoint and the Call Agent controlled T.38 fax relay procedure was initiated. The Call Agent SHOULD modify each side of the connection to start using the "image/t38" media format, unless they already do. Note that, as long as use of the Call Agent controlled T.38 relay procedure is in effect, the event will be generated upon fax call detection irrespective of the current encoding method on any connections on the endpoint (incl. "image/t38"). The T38(start) event MUST be generated at most once by the endpoint per fax call regardless of whether or not it is requested again in a subsequent requested events list.

* stop Call Agent controlled T.38 fax relay procedure ended and the gateway did not detect any errors. Note that this does not necessarily imply a successfully transmitted fax. It merely indicates that the Call Agent controlled T.38 fax relay procedure has ended and the procedure itself did not encounter any errors. The Call Agent may want to modify the media parameters for each side of the connection. Note that, in contrast to the gateway controlled fax procedure case, media parameters such as codecs do not automatically revert to their values before the start of the fax call; echo cancellation and silence suppression however do per the procedures in [RFC3435] Section 2.3.5. The "t38(stop)" event MUST NOT be generated unless a corresponding "t38(start)" event for the fax call in question was generated earlier.

* failure Call Agent controlled T.38 fax relay procedure ended abnormally. Some kind of problem in the Call Agent controlled T.38 fax relay procedure was encountered and the procedure ended. The Call Agent may want to modify the media parameters for each side of the connection. Note that, in contrast to the gateway controlled fax procedure case, media parameters such as codecs do not automatically revert to their state before the start of the fax call; echo cancellation and silence suppression however do per the procedures in [RFC3435] Section 2.3.5. The "t38(failure)" event MUST NOT be generated unless a corresponding "t38(start)" event for the fax call in question was generated earlier.

One of the above parameters will be present when the event is reported. The "t38" event MAY be parameterized with additional parameters, however it is RECOMMENDED that one of the above parameters is the first parameter supplied. Unknown parameters MUST be ignored.

The following example illustrates the encoding of the "t38" event:

```
O: fxr/t38(start)
O: fxr/t38(fooobar, stop)
```

2.3 Connection Parameters

The connection parameters for the connection, that measures packets and octets sent and received, MUST include packets and octets for fax handling as well. Interarrival jitter and average transmission delay calculation however MAY not be performed while fax is in progress, e.g., if T.38 is used. In such cases, the interarrival jitter and average transmission delay calculations are simply suspended until calculations can resume, e.g., by changing back to an RTP-based media stream again.

In addition to these connection parameters, the fax package defines the following connection parameters, which gateways MAY support:

Number of fax pages sent (PGS):

The cumulative number of fax pages sent by the endpoint for the life of the connection. The parameter is encoded as "PGS" and the value supplied is a string of up to nine decimal digits.

Number of fax pages received (PGR):

The cumulative number of fax pages received by the endpoint for the life of the connection. The parameter is encoded as "PGR" and the value supplied is a string of up to nine decimal digits.

The following example illustrates the use of these parameters:

P: FXR/PGS=3, FXR/PGS=0, PS=1245, OS=62345, ...

2.4 Negotiation of T.38 Parameters

T.38 Annex D defines a number of T.38 parameters that can be negotiated in the SDP. Currently, T.38 does not specify procedures for how each of these parameters is negotiated, and in particular whether each side has to use the same value or not. In the absence of that, it was considered to add such definitions and procedures here. However, it is expected that T.38 will rectify the above, which could lead to conflicting definitions and procedures. To avoid that, we instead assume the existence of an offer/answer section for T.38 where T.38 Annex D parameters are classified as either declarative or negotiated, and we then provide guidelines for how to map such definitions and procedures to the MGCP fax package defined here.

MGCP does not specify use of the offer/answer model, but instead operates with the concept of connection handling commands (e.g., CreateConnection and ModifyConnection) that may include a RemoteConnectionDescriptor (SDP) and in turn may generate a LocalConnectionDescriptor (SDP) in their response.

When an MGCP endpoint receives a CreateConnection command without a RemoteConnectionDescriptor, it should follow the corresponding T.38 procedures for generating an initial offer and return the resulting the SDP in its LocalConnectionDescriptor.

When an MGCP endpoint receives a CreateConnection command with a RemoteConnectionDescriptor, it should follow the corresponding T.38 procedures for receiving an initial offer and generating an answer to it. The resulting SDP is returned in the LocalConnectionDescriptor.

When an MGCP endpoint receives a ModifyConnection command with a RemoteConnectionDescriptor, it cannot determine whether this corresponds to an answer to an initial offer or a new offer. This is not an issue for declarative parameters since they can be specified independently in either direction. Negotiated parameters however require some consideration:

When an offerer receives an answer to a previous offer, the negotiation has completed and the parameters negotiated can no longer be changed with this offer/answer exchange. The negotiated parameters may be subject to certain validation checks. Conversely, when an answerer receives an offer, the negotiation is open and the answerer may change some of the offered negotiated parameters. Since the MGCP endpoint does not know which situation it is in, it can not perform the "offerer" validation checks. Likewise, in order to ensure that any required negotiation actually takes place, it needs to process incoming SDP as an offer. If the SDP in fact does correspond to an offer then this is obviously correct behavior. If the SDP however corresponds to an answer, and one or more negotiated parameters did change, then this will result in a new SDP. The Call Agent may or may not contain sufficient intelligence to determine whether this new SDP needs to result in another offer/answer exchange or not.

For example, if the initial offer (from CreateConnection without SDP) contained fax version 2, and the answer (from CreateConnection with SDP) contained fax version 0, then the corresponding ModifyConnection command (with SDP) will result in an updated

SDP with fax version also set to zero. If this was the only change in the updated SDP, a new offer/answer exchange would not be needed. Note that this example does not imply that it is generally considered a good idea for Call Agents to parse SDP in order to determine whether new offer/answer exchanges are needed or not.

Finally, a ModifyConnection without SDP that generates an SDP needs to be considered. The SDP generated may either correspond to an initial offer/answer exchange or a subsequent offer/answer exchange. The endpoint should generate SDP as if it was part of a subsequent offer/answer exchange. If the Call Agent does not desire such semantics, it can simply create a new connection instead.

2.5 Implementation Considerations

2.5.1 Media IP Address and Port for T.38

When an endpoint is instructed to change to or from T.38 for a media stream, it SHOULD continue using the same IP address and port as the media stream is currently using, since this will minimize any Quality of Service, Network Address Translator (NAT) and Firewall interactions from the change. However, if an endpoint has a good reason, it MAY choose not to follow this recommendation.

When an endpoint uses the same port for RTP audio and T.38 with either UDPTL or TCP, packets of one type (e.g., T.38) may be received while expecting packets of another type (RTP audio). Since there is explicit signaling to indicate which type is expected at any given point in time, this does not introduce any new problems. In other words, the receiver does not operate as a demultiplexer with a need to determine if a given packet received is an RTP audio packet or a T.38 UDPTL/TCP packet. The receiver simply processes incoming packets as usual. If T.38 packets are expected, then incoming packets are validated against T.38, and if RTP audio packets are expected, then incoming packets are validated against RTP.

2.5.2 Case Sensitivity

The IANA has registered the uppercase string "UDPTL" as the transport protocol identifier to be used for UDP-based T.38. However, the examples provided in Recommendation T.38 as well as most (if not all) current implementations use the lowercase string "udptl" instead. Implementations conforming to this package SHOULD generate the lowercase string "udptl" and accept the lowercase, uppercase, and mixed upper/lowercase strings as being equivalent.

The attribute "T38MaxBitRate" is incorrectly registered with IANA as "T38maxBitRate" (lower-case "m"). In accordance with T.38 examples and common implementation practice, the form "T38MaxBitRate" SHOULD be generated by implementations conforming to this package.

In general, it is RECOMMENDED that implementations of this package accept both lowercase, uppercase, and mixed upper/lowercase encodings of all the T.38 attributes.

2.5.3 Boolean Indicator after T.38 Parameters

Some implementations incorrectly use a colon (':') followed by a number (zero or one) after the attributes T38FaxFillBitRemoval, T38FaxTranscodingMMR and T38FaxTranscodingJBIG. Implementations that receive such erroneous encodings MAY interpret the value ":0" as lack of support for the option and all other values as indicating support of the option in question.

3. Call Flow Examples

In this section, we provide two example call flows. The first one illustrates a T.38 fax call under Call Agent control on both the originating and terminating side. The second one illustrates the use of multiple and different options on the two sides.

3.1 Call Agent Controlled T.38 Strict

In this example, both sides are under strict T.38 Call Agent control. We assume the originating and terminating Call Agent communicate via the Session Initiation Protocol (SIP) [RFC3261] (also, see [SIPfax]):

#	GW-o	CA-o	CA-t	GW-t
1	<- CRCX			
2	200(sdp-o) ->			
3		INVITE(sdp-o) ->		
4			CRCX(sdp-o) ->	
5			<- 200 (sdp-t)	
6		<- 200(sdp-t)		
7	<- MDCX(sdp-t)			
8	200 ->			
9				<- ANS/ T.30 CED
10				<- T.30 fax preamble
11				<- NTFY(t38 start)
12			200 ->	
13			MDCX(t38) ->	
14			<- 200(sdp-t2)	
15		<- INVITE(sdp-t2)		
16	<- MDCX(sdp-t2)			
17	200(sdp-o2) ->			
18		200(sdp-o2) ->		
19			MDCX(sdp-o2) ->	
20			<- 200	
21				(fax ends)
22			<- NTFY(t38 stop)	
24			200 ->	

Step 1:

The Call Agent issues a CreateConnection command to the gateway instructing it to use PCMU media encoding and to use the strict Call Agent controlled T.38 procedure. Consequently, the Call Agent asks the gateway to notify it of the t38 event:

```
CRCX 1000 ds/ds1-1/1@gw-o.whatever.net MGCP 1.0
C: 1
L: a:PCMU, fxr/fx:t38
M: recvonly
R: fxr/t38
X: 1
```

Step 2:

The gateway acknowledges the command and includes SDP with codec information as well as RFC 3407 capability information:

```
200 1000 OK
I:1

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
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 3:

The originating Call Agent sends a SIP INVITE message with the SDP to the terminating Call Agent.

Step 4:

The terminating Call Agent issues a CreateConnection command to the terminating gateway instructing it to use PCMU media encoding and to use the strict Call Agent controlled T.38 procedure. Consequently, the Call Agent asks the gateway to notify it of the t38 event:

```
CRCX 2000 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
L: a:PCMU, fxr/fx:t38
M: sendrecv
R: fxr/t38
X: 20

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
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 5:

The terminating gateway supports T.38, and the RemoteConnectionDescriptor included indicates that the other side supports T.38 as well, so the strict T.38 Call Agent controlled procedure requested can be used. The terminating gateway sends back a success response with its SDP which also includes capability information:

```

200 2000 OK
I:2

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38

```

Step 6:

The terminating Call Agent sends back a SIP 200 OK response to the originating Call Agent, which in turn sends a SIP ACK (not shown).

Step 7:

The originating Call Agent in turns sends a ModifyConnection command to the originating gateway:

```

MDCX 1001 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
M: sendrecv

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38

```

The ModifyConnection command does not repeat the LocalConnectionOptions sent previously. As far as fax handling is concerned, the gateway therefore attempts to continue using the current fax handling procedure, i.e., strict Call Agent controlled T.38. Since the capability information indicates the other side supports T.38, the gateway will in fact be able to use the strict Call Agent controlled T.38 procedure. Had there not been any support for T.38 in the RemoteConnectionDescriptor, then this command would still have succeeded, however there would be no special fax handling procedure (since strict mode could not be supported).

Step 8:

The gateway acknowledges the command. At this point, a call is established using PCMU encoding, and if a fax call is detected, the Call Agent controlled T.38 procedure will be initiated.

Step 9-11:

A fax call now occurs. First, the T.30 CED tone (a.k.a. V.25 ANS) is sent which in this case is simply passed through the current PCMU encoding. Since both fax and modem calls can start with this sequence, it is not possible to determine that this is a fax call until step 10, where the V.21 fax preamble is detected.

The gateway was instructed to apply the Call Agent controlled T.38 procedure for fax calls, so the "t38(start)" event occurs and is notified to the Call Agent:

```
NTFY 2500 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: fxr/t38(start)
X: 20
```

Step 12:

The Call Agent acknowledges the Notify command:

```
200 2500 OK
```

Step 13:

The Call Agent then instructs the terminating gateway to change to using the "image/t38" MIME type instead:

```
MDCX 2002 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2
L: a:image/t38
R: fxr/t38
X: 21
```

Step 14:

The gateway changes to T.38, and sends back a success response with updated SDP:

```
200 2002 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udptl t38
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Note, that since the gateway's current RemoteConnectionDescriptor (as opposed to the LocalConnectionDescriptor returned here) does not list "image/t38" as a valid encoding method, the terminating gateway is still muting the media and is now waiting for an updated RemoteConnectionDescriptor with "image/t38".

Step 15:

The terminating Call Agent sends a re-INVITE to the originating Call Agent with the updated SDP.

Step 16:

The originating Call Agent then sends a ModifyConnection command to the originating gateway:

```
MDCX 1003 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 17:

The originating gateway changes to T.38 and sends back a success response with updated SDP:

```
200 1003 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 18:

The originating Call Agent sends a SIP 200 OK response with the updated SDP to the terminating Call Agent, which in turn sends a SIP ACK (not shown).

Step 19:

The terminating Call Agent sends a ModifyConnection with the updated SDP to the terminating gateway:

```
MDCX 2003 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2
```



```

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38

```

Step 20:

The terminating gateway sends back a success response:

```
200 2003 OK
```

Since the terminating gateway now has a RemoteConnectionDescriptor with "image/t38" as valid media, it can start exchanging T.38 with the originating gateway.

Step 21, 22:

When the fax ends, a "t38(stop)" event is generated, which is notified to the Call Agent:

```

NTFY 2501 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: t38(stop)
X: 3

```

Step 23:

The Call Agent acknowledges the Notify command:

```
200 2501 OK
```

The fax call is now over. The Call Agent may now decide to change back to a voice codec, delete the connection, or something different.

3.2 Multiple and Different Options

In this example, the originating gateway is instructed to use the gateway procedure whereas the terminating gateway is given a choice between gateway procedure and strict t38 procedure. Furthermore, the originating fax machine is generating CNG tone.

#	GW-o	CA-o	CA-t	GW-t
1	<-	CRCX		
2	200(sdp-o) ->			
3		INVITE(sdp-o) ->		
4			CRCX(sdp-o) ->	
5			<-	200 (sdp-t)
6		<-	200(sdp-t)	
7	<-	MDCX(sdp-t)		
8	200 ->			

9	CNG ->			
10				<- ANS/T.30 CED
11				<- T.30 fax p.
12				<- NTFY(t38 start)
13			200 ->	
14			MDCX(t38) ->	
15			<- 200(sdp-t2)	
16		<- INVITE(sdp-t2)		
17	<- MDCX(sdp-t2)			
18	200(sdp-o2) ->			
19		200(sdp-o2) ->		
20		MDCX(sdp-o2) ->		
21		<- 200		
22				(fax ends)
23				<- NTFY(t38 stop)
24			200 ->	

Step 1:

The Call Agent issues a CreateConnection command to the gateway instructing it to use PCMU media encoding and to use the gateway procedure. Consequently, the Call Agent asks the gateway to notify it of the gwfax event:

```
CRCX 1000 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
L: a:PCMU, fxr/fx:gw
M: recvonly
R: fxr/gwfax
X: 1
```

Step 2:

The gateway acknowledges the command and includes SDP with codec information as well as capability information:

```
200 1000 OK
I:1

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
a=sqn: 0
a=cddsc: 1 audio RTP/AVP 0 18
a=cddsc: 3 image udptl t38
a=X-FaxScheme123
```

We assume the gateway supports some other fax scheme and it indicates this by including an attribute "FaxScheme123"

Step 3:

The originating Call Agent sends a SIP INVITE message with the SDP to the terminating Call Agent.

Step 4:

The terminating Call Agent issues a CreateConnection command to the terminating gateway instructing it to use PCMU media encoding and to use either the gateway procedure or strict Call Agent controlled T.38 procedure. Consequently, the Call Agent asks the gateway to notify it of both the gwfax and t38 events:

```
CRCX 2000 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
L: a:PCMU, fxr/fx:gw;t38
M: sendrecv
R: fxr/t38, fxr/gwfax
X: 20
```

```
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
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
a=X-FaxScheme123
```

Step 5:

The terminating gateway does not support any special gateway fax handling, however it does support T.38, and the RemoteConnectionDescriptor included indicates that the other side supports T.38 as well, so the strict T.38 Call Agent controlled procedure requested can be honored. The terminating gateway sends back a success response with its SDP which also includes capability information:

```
200 2000 OK
I:2

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 6:

The terminating Call Agent sends back a SIP 200 OK response to the originating Call Agent, which in turn sends a SIP ACK (not shown).

Step 7:

The originating Call Agent in turns sends a ModifyConnection command to the originating gateway:

```
MDCX 1001 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
M: sendrecv

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

The ModifyConnection command does not repeat the LocalConnectionOptions sent previously. As far as fax handling is concerned, the gateway therefore attempts to continue using the current fax handling, i.e., the gateway procedure. The SDP information returned however does not indicate support for the "FaxScheme123", and hence the originating gateway will not invoke any special fax handling procedure for this call.

Step 8:

The gateway acknowledges the command. At this point, a call is established using PCMU encoding and if a fax call is detected, no special fax handling procedure will occur.

Step 9-12:

First, a CNG tone is generated by the originating fax thereby indicating a fax call. If the gateway was using either of the T.38 modes, or it had negotiated support for special gateway handling procedure with the other side, a "t38(start)" or "gwfax(start)" event would now have been generated and the switch to T.38 (or special gateway handling) could start. However, since the negotiation with the terminating gateway resulted in the originating gateway not doing anything special for fax, no such event is generated. Instead, the "nopfax(start)" event is now generated, but since the Call Agent has not requested this event, it is not detected and hence not reported to the Call Agent. Consequently, the CNG tone is simply passed through the current PCMU encoding without the (originating) Call Agent being aware of the fax call.

Subsequently, the T.30 CED tone (a.k.a. V.25 ANS) occurs which in this case is also simply passed through the current PCMU encoding. Since both fax and modem calls can start with this sequence, it is not possible to determine that this is a fax call until step 11, where the V.21 fax preamble is detected.

The terminating gateway is using the Call Agent controlled T.38 procedure for fax calls, so the "t38(start)" event occurs and is notified to the Call Agent:

```
NTFY 2500 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: fxr/t38(start)
X: 20
```

Step 13:

The Call Agent acknowledges the Notify command:

```
200 2500 OK
```

Step 14:

The Call Agent then instructs the terminating gateway to change to using the "image/t38" MIME type instead:

```
MDCX 2002 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2
L: a:image/t38
R: fxr/t38
X: 21
```

Step 15:

The gateway changes to T.38, and sends back a success response with updated SDP:

```
200 2002 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udpt1 t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udpt1 t38
```

Note, that since the terminating gateway's last received RemoteConnectionDescriptor (as opposed to the LocalConnectionDescriptor returned here) did not list "image/t38" as a valid encoding method, the terminating gateway is still muting the media and is now waiting for an updated RemoteConnectionDescriptor with "image/t38".

Step 16:

The terminating Call Agent sends a re-INVITE to the originating Call Agent with the updated SDP.

Step 17:

The originating Call Agent then sends a ModifyConnection command to the originating gateway:

```
MDCX 1003 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udptl t38
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 18:

The originating gateway changes to T.38 and sends back a success response with updated SDP:

```
200 1003 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 19:

The originating Call Agent sends a SIP 200 OK response with the updated SDP to the terminating Call Agent, which in turn sends a SIP ACK (not shown).

Step 20:

The terminating Call Agent sends a ModifyConnection with the updated SDP to the terminating gateway:

```
MDCX 2003 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
```

```

m=image 3456 udptl t38
a=sqn: 0
a=csrc: 1 audio RTP/AVP 0 18
a=csrc: 3 image udptl t38

```

Step 21:

The terminating gateway sends back a success response:

```
200 2003 OK
```

Since the terminating gateway now has a RemoteConnectionDescriptor with "image/t38" as valid media, it can start exchanging T.38 with the originating gateway.

Step 22, 23:

When the fax ends, a "t38(stop)" event is generated, which is notified to the Call Agent:

```

NTFY 2501 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: t38(stop)
X: 3

```

Step 24:

The Call Agent acknowledges the Notify command:

```
200 2501 OK
```

The fax call is now over. The Call Agent may now decide to change back to a voice codec, delete the connection, or something different.

3.3 Interaction with SIP Endpoints

In this example, we show interaction with a SIP endpoint that does not support the RFC 3407 capability descriptors. To accommodate such endpoints, the T.38 loose mode is being used (at the risk of initiating T.38 to an endpoint that does not support it).

#	GW-o	CA-o	SIP-UA-t	fax
1	<-	CRCX		
2	200(sdp-o)	->		
3		INVITE(sdp-o)	->	
4		<-	200(sdp-t)	
5		ACK	->	
6	<-	MDCX(sdp-t)		
7	200	->		
8				<- ANS/ T.30 CED
9				<- T.30 fax preamble
10		<-	INVITE(sdp-t2)	
11	<-	MDCX(sdp-t2)		

12	200(sdp-o2) ->		
13		200(sdp-o2) ->	
14		<- ACK	
--	-----	-----	-----
15			(fax ends)
16		<- BYE	
17		200 ->	

Step 1:

The Call Agent issues a CreateConnection command to the gateway instructing it to use PCMU media encoding and to use the loose Call Agent controlled T.38 procedure. Consequently, the Call Agent asks the gateway to notify it of the t38 event:

```
CRCX 1000 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
L: a:PCMU, fxr/fx:t38-loose
M: recvonly
R: fxr/t38
X: 1
```

Step 2:

The gateway acknowledges the command and includes SDP with codec information as well as RFC 3407 capability information:

```
200 1000 OK
I:1

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
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 3:

The originating Call Agent sends a SIP INVITE message with the SDP to the terminating Call Agent (not all SIP details shown here):

```
INVITE sip:bob@biloxi.example.com SIP/2.0
...
Content-Type: application/sdp
Content-Length: 167

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
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 4:

The terminating SIP User Agent sends back a SIP 200 OK response (not all SIP details shown) to the originating Call Agent:

```
SIP/2.0 200 OK
...
Content-Type: application/sdp
Content-Length: 100

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
```

Note that the terminating SIP User Agent does not use the RFC 3407 capability descriptor to indicate support for (or lack of support for) T.38.

Step 5:

The originating Call Agent receives the SIP 200 response and sends a SIP ACK message to the terminating SIP UA.

Note that the Call Agent does not know whether the peer entity supports T.38 or not. In order to figure this out, the Call Agent could send a SIP OPTIONS request to the terminating SIP UA requesting it to return its capabilities (not shown). Note that this can of course be done towards any SIP peer, e.g., if the other side was a Call Agent speaking SIP it could be done there too.

Step 6:

The originating Call Agent in turns sends a ModifyConnection command to the originating gateway:

```
MDCX 1001 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
M: sendrecv

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 0
```

The ModifyConnection command does not repeat the LocalConnectionOptions sent previously. As far as fax handling is concerned, the gateway therefore attempts to continue using the current fax handling procedure, i.e., loose Call Agent controlled T.38. The T.38 loose procedure can always be supported, and hence a switch to T.38 will be attempted if the originating gateway detects a fax call.

Step 7:

The gateway acknowledges the command. At this point, a call is established using PCMU encoding, and if a fax call is detected, the Call Agent controlled T.38 procedure will be initiated.

Step 8-9:

A fax call now occurs. First, the T.30 CED tone (a.k.a. V.25 ANS) is sent which in this case is simply passed through the current PCMU encoding. Since both fax and modem calls can start with this sequence, it is not possible to determine that this is a fax call until step 9, where the V.21 fax preamble is detected.

Step 10:

The terminating SIP UA in fact does support T.38, and upon detecting the fax call, it attempts to change to T.38. Consequently, it sends a re-INVITE to the originating Call Agent with an updated SDP indicating a switch to T.38.

```

INVITE sip:ca@ca-o.whatever.net SIP/2.0
...
Content-Type: application/sdp
Content-Length: 100

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udpt1 t38

```

Step 11:

The originating Call Agent then sends a ModifyConnection command to the originating gateway:

```

MDCX 1003 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udpt1 t38

```

Step 12:

The originating gateway changes to T.38 and sends back a success response with updated SDP:

```
200 1003 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 13:

The originating Call Agent sends a SIP 200 OK response with the updated SDP to the terminating SIP User Agent:

```
SIP/2.0 200 OK
...
Content-Type: application/sdp
Content-Length: 167

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 0 18
a=cdsc: 3 image udptl t38
```

Step 14:

The terminating SIP User Agent receives the SIP 200 and sends a SIP ACK.

Since the terminating SIP User Agent now has a RemoteConnectionDescriptor with "image/t38" as valid media, it can start exchanging T.38 with the originating gateway (and vice versa).

Step 15-17:

When the fax ends, the terminating SIP UA decides to tear down the call and hence sends a SIP BYE message, which the Call Agent responds to with a SIP 200.

The Call Agent may now decide to change back to a voice codec, delete the connection, or something different.

4. Security Considerations

The MGCP fax package itself is not known to introduce any new security concerns. However, implementers should note, that T.38 media is currently transported over UDP (UDPTL) or TCP in clear and without any integrity protection. If for example security services are in place to protect RTP media streams, these will thus not be in effect for the T.38 media stream. If such lack of security is a concern, the fax LocalConnectionOptions allowing T.38 in this package SHOULD NOT be used, i.e., the "off" (or a new secure extension) fax LocalConnectionOption should be used.

5. IANA Considerations

The IANA is hereby requested to register the following MGCP package:

Package Title	Name	Version
-----	----	-----
Fax	FXR	0

6. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3435] F. Andreassen, B. Foster, "Media Gateway Control Protocol (MGCP) Version 1.0", RFC 3435, January 2003.
- [T38] ITU-T Recommendation T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", 03/2002.
- [RFC3407] F. Andreassen, "Session Description Protocol (SDP) Simple Capability Declaration", RFC 3407, October 2002.

7. Informative References

- [T30] ITU-T Recommendation T.30, "Procedures for document facsimile transmission in the general switched telephone network", 07/03.
- [RFC3261] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [SIPfax] Mule, J., and J. Li, "SIP Support for Real-time Fax: Call Flow Examples and Best Current Practices", work in progress.

8. Acknowledgements

Several people have contributed to the development of the MGCP fax package. In particular, the author would like to thank Bill Foster, Paul Jones, Gary Kelly, Rajesh Kumar, Dave Horwitz, Hiroshi Tamura, Rob Thompson and the CableLabs PacketCable NCS focus team for their contributions.

9. Author's Address

Flemming Andreassen
Cisco Systems

499 Thornall Street, 8th Floor
Edison, NJ 08837

Email: fandreas@cisco.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the IETF's procedures with respect to rights in IETF Documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

A.4 VoIP Metrics Package

The VoIP Metrics package provides a mechanism to control the collection and reporting of VoIP metrics. The MTA MUST support the reporting of VoIP Metric parameters as specified in this package for all Extended Report parameters mandated in [11]. The Implementations MUST conform to the NCS Protocol Version encoding as specified in Section 8.2.1.4.

Internet Engineering Task Force

D. Auerbach
Cisco Systems
D. Hancock

CableLabs
B. Hare
Arris
R. Kumar
Cisco Systems

Internet Draft
Document:.txt
Category: Informational

October 1, 2004

Media Gateway Control Protocol VoIP RTCP-XR Metrics Package

Status of this Memo

By submitting this Internet-Draft, we certify that any applicable patent or other IPR claims of which we are aware have been disclosed, and any of which we become aware will be disclosed, in accordance with RFC 3668.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 1, 2005.

Abstract

This document defines a Media Gateway Control Protocol (MGCP) package to control the collection of metrics supported by RTCP Extended Report as specified in RFC 3611. The package allows for the call agent to enable or disable the collection of these metrics. It also allows the call agent to control whether or not the gateway will respond to remote requests for metric collection from a peer gateway.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [RFC2119].

Table of Contents

- 1. Introduction
- 2. VoIP Metrics Package Definition
 - 2.1 LocalConnectionOptions
 - 2.1.1 Endpoint Procedures
 - 2.1.1.1 Off Procedure
 - 2.1.1.2 On Procedure
 - 2.1.1.3 Negotiate Procedure
 - 2.2 AuditConnection procedures
 - 2.3 New MGCP Extension Parameters
 - 2.3.1 Extension Parameter Values
 - 2.3.2 Extension Parameter Grammar
 - 2.4 VoIP Metrics Negotiation
 - 2.5 Implementation Considerations
 - 2.5.1 Inter-working Considerations
- 3. Call Flow Examples
 - 3.1 Call agent initiated DLCX
 - 3.2 Gateway initiated DLCX
- 4. Security Considerations
- 5. IANA Considerations
- 6. Normative References
- 7. Acknowledgments

1. Introduction

This document defines a Media Gateway Control Protocol (MGCP) [RFC3435] package that enables the collection and reporting of metrics that measure the quality of voice traffic on VoIP connections. This capability is based on and utilizes the VoIP Metrics Block of the RTCP Extended Report as specified in RFC 3611.

In addition to the parameters in the RTCP-XR VoIP metrics block, this document includes the inter-arrival jitter as defined in RFC 3550. The purpose of this inclusion is the alignment of MGCP-based VoIP metrics reporting with SIP. Note that the MGCP ConnectionParameters structure [RFC3435] allows the reporting of the local, but not the remote, version of this parameter.

This package makes a distinction between collecting, responding and reporting. The term "collecting" refers to the process of calculating and accumulating local VoIP metrics data and accumulating remote VoIP metrics data received in RTCP XR packets from the remote device. The term "responding" is the sending of the accumulated local metrics data in RTCP XR packets to the remote connected device as specified in RFC 3611. The term "reporting" refers to the process of reporting the accumulated local and remote VoIP metric data for the connection to the Call Agent in either a DeleteConnection or AuditConnection command sequence.

The Call Agent can ask a local endpoint to enable (or disable) VoIP metrics reporting for a connection in the CreateConnection or ModifyConnection command. When VoIP metrics reporting is enabled, the local endpoint will use SDP signaling as specified in RFC 3611 to ask the remote endpoint to start transmitting RTCP XR VoIP metrics data for the remote end of the connection. The local endpoint will accumulate a local and remote copy of the VoIP metrics data to be reported to the Call Agent when the connection is deleted, or audited (and the metrics data is requested).

An endpoint will collect local VoIP metric data for two separate reasons; when VoIP metrics reporting is enabled by the Call Agent via connection command, or when requested to collect VoIP metrics by the remote endpoint via SDP signaling. The Call Agent can inhibit the collection and responding to remote endpoints by specifically turning off the collection/responding function. The VoIP metrics data that is reported to the Call Agent when the connection is deleted represents the data that was accumulated since the endpoint started collecting originally on the connection, whether collecting was initiated by a Call Agent request or by SDP signaling.

This package is designed to provide an accurate measure of voice quality for basic 2-way call connections, where metrics are accumulated over the life of the connection and reported to the Call Agent when the connection is deleted. The package does not make any special provisions for mid-call path reconfiguration events that could affect voice quality. For example, if a ModifyConnection command changes the remote endpoint of a connection, the single metrics report generated when the connection is subsequently deleted would not provide any means to compare the metrics before and after the remote endpoint was changed.

Note that this package does not affect the reporting of other MGCP statistics as defined in the base protocols.

The VoIP Metrics package definition is provided in Section 2 and in Section 3 we provide two call flow examples showing how to use it. Security considerations are found in Section 4, followed by the IANA considerations and references.

2. VoIP Metrics Package Definition

A package is defined for VoIP Metrics. The package defines new localconnectionoptions, and connection parameters as detailed below.

Package Name: XRM

Package Version: 0

2.1 LocalConnectionOptions

This package defines a new LocalConnectionOptions (LCO) parameter to enable and disable VoIP Metrics reporting at the gateway. This is an optional parameter with a default value of "negotiate". The Call Agent supplies this LCO parameter to instruct the endpoint to enable or disable the collection and reporting of RTCP-XR Metrics for the target connection. The new parameter is the metric collection (and) report parameter and is encoded as "mcr".

The parameter can have one of the following three values:

- off - disable VoIP metrics collecting, reporting and responding (encoded as "off").

- on - enable VoIP metrics collecting, reporting and responding (encoded as "on").
- negotiate - allow VoIP metrics collecting and responding if requested but disable VoIP reporting (encoded as "negotiate").

The following example shows how to use this parameter to turn VoIP metrics collection, reporting and responding on:

```
L:      xrm/mcr:on
```

In CreateConnection commands, the xrm/mcr LCO value defaults to "negotiate". In ModifyConnection commands, the xrm/mcr LCO value defaults to its current value for the connection. Thus, if LocalConnectionOptions are either omitted or the xrm/mcr LCO is not included in a ModifyConnection command, the previous xrm/mcr LCO value for the connection is retained.

Note that this LCO directly controls the reporting of XR metrics to the Call Agent and may control the collection of metrics. The collection of metrics is also dependent on a possible request from a peer gateway as indicated in the RemoteConnectionDescriptor.

2.1.1 Endpoint Procedures

This section describes the endpoint procedures for each of the xrm/mcr parameter values.

2.1.1.1 Off Procedure

On receiving this value, the endpoint MUST erase any local or remote VoIP metric data previously collected for the connection, and MUST instruct the far-end device to stop sending RTCP XR VoIP Metrics Report Blocks. The endpoint MUST ignore an SDP request for metrics and MUST NOT transmit local metrics to the peer endpoint. The local endpoint MUST NOT report XR metrics upon receiving or initiating a DeleteConnection sequence and MUST NOT report XR metrics when responding to an AuditConnection command.

2.1.1.2 On Procedure

On receiving this value, if not already doing so, the endpoint MUST start calculating local VoIP Metrics data for the connection, and MUST instruct the far-end device to start sending VoIP Metrics data for the connection by including an "rtcp-xr" attribute with a value of "voip-metrics" in the SDP.

The endpoint MUST report local and remote VoIP metrics in the response to DeleteConnection, in the response to AuditConnection when the metric parameters are audited, and when sending a DeleteConnection command. The endpoint MUST omit the remote metrics from the report if no remote metrics were received from the remote endpoint.

2.1.1.3 Negotiate Procedure

On receiving this value, the endpoint does nothing except to set an internal flag to indicate that SDP request for local metrics to be sent to a remote peer endpoint are allowed. The local does not send this information until it is requested to do so by the remote endpoint.

2.2 AuditConnection procedures

This package defines two new RequestedInfo values to be used with an AuditConnection (AUCX). These allow a call agent to explicitly request that the endpoint return the current values of either the local or remote metrics. The request does not reset the values. When one or both of these request info values are included in an AUCX command, the endpoint MUST return the appropriate metric values if available. The two new request info values are:

XRM/LVM	Return the local voice metric values.
XRM/RVM	Return the remote voice metric values.

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 audit response.

If an endpoint is queried about a parameter it does support, but has no value for, the endpoint MUST NOT generate an error; instead the parameter MUST be included in the audit response with an empty parameter value.

2.3 New MGCP Extension Parameters

Two new MGCP parameters are defined to support VoIP metrics reporting; one for reporting local metrics data, and one for reporting remote metrics data. They are returned in the DeleteConnection sequence and may be returned by the AuditConnection response (when the audit is requested). Local metrics are included in the "XRM/LVM" line. Remote metrics are included in the XRM/RVM line.

As an example, for a gateway that is recording local metrics and has received remote metrics, a DLCX will result in the following additional lines being returned in the response:

```
XRM/LVM: PLC=1, JBA=2, JBR=7 ...  
XRM/RVM: PLC=0, JBA=0, JBR=3 ...
```

For the detailed semantics of these parameters, see the description of the VoIP Metrics Report block in RFC 3611. For example, the meaning of "standard" vs. "enhanced" packet loss concealment must be gleaned from RFC 3611.

2.3.1 Extension Parameter Values

Unlike RTCP, parameter fields in MGCP are expressed in a text rather than a fixed binary format. Therefore, it is not necessary to stipulate upper bounds for numeric parameters in MGCP. However, the upper bounds specified in RFC 3611 still apply. Note that upper bounds in RFC 3611 are sometimes dictated by regularity in field size (e.g., increments of 8) rather than by the actual range of parameter values. For example, although a numeric value of 65,535 ms can be expressed as a 16-bit field, a viable RTP connection will never have a round trip delay with this value.

Extension Parameter Name	Code	Value
Network packet loss rate	NLR	The proportion of packets lost since the start of transmission expressed as an 8-bit binary fraction obtained by dividing the number of packets lost in the transmission path by the total number of packets expected by the receiver, multiplying this value by 256 and taking the integer portion of the result.
Jitter buffer discard rate	JDR	The proportion of the packets discarded by the receiving jitter buffer since the start of transmission expressed as an 8-bit binary fraction obtained by dividing the number of packets discarded by the jitter buffer by the total number of packets expected by the receiver, multiplying this value by 256 and taking the integer portion of the result.
Burst loss density	BLD	The average proportion of lost and discarded packets occurring during burst periods expressed as an 8 bit binary fraction obtained by dividing the sum of the number of packets detected as lost during burst periods and the number of packets discarded by the jitter buffer during burst periods by the total number of packets expected by the receiver, multiplying this value by 256 and taking the integer portion of the result.
Gap loss density	GLD	The average proportion of lost and discarded packets occurring during gap periods expressed as an 8 bit binary fraction obtained by dividing the sum of the number of packets detected as lost during gap periods and the number of packets discarded by the jitter buffer during gap periods by the total number of packets expected by the receiver, multiplying this value by 256 and taking the integer portion of the result.
Burst duration	BD	The average duration of the burst periods, in milliseconds.
Gap duration	GD	The average duration of the gap periods, in milliseconds.
Round trip network delay	RTD	The round trip delay, in milliseconds, between the RTP interfaces of the local and remote call endpoints. Valid values MUST be greater than or equal to 0 milliseconds.
End system delay	ESD	The end system (endpoint) delay, in milliseconds, comprising the encode, decode and jitter buffer delay. Valid values MUST be greater than or equal to 0 milliseconds.
Signal level	SL	The ratio of the average signal level to a 0 dBm0 reference, expressed in dB. This is expressed as an integer with an optional negative ("-") sign. When the sign is omitted, a positive value is assumed. Typical values should generally be between -15 and -20.
Noise level	NL	The ratio of the silent period average background noise to a 0 dBm0 reference, in dB. No sign is necessary since valid values MUST be less than or equal to 0 dB.
Residual echo return loss	ERL	The echo return loss after the effects of echo cancellation, in dB. No sign is necessary since valid values MUST be greater than or equal to 0 dB.
Minimum gap threshold	GMN	The gap/burst transition threshold; the recommended value is 16. Values in the range 1 - 255 are permitted.
R factor	NSR	A value representing the receiving end call quality for the RTP packet stream originated from the reporting endpoint; calculated as per ITU-T Recommendation G.107. Valid values are greater than 0 and less than or equal to 100.

External R factor	XSR	A value representing the effects of any call segment carried over a network segment external to the one on which the endpoint resides; calculated as per ITU-T Recommendation G.107. Valid values are greater than 0 and less than or equal to 100.
Estimated MOS-LQ	MLQ	An estimated receiving end Listening Quality MOS. The nominal range of MOS scores is 0-5. Before being expressed in MGCP, the MOS score is multiplied by 10 and any fractional part is truncated. This results in an integer with valid values in the range 10 -50.
Estimated MOS-CQ	MCQ	An estimated receiving end Conversational Quality MOS. The nominal range of MOS scores is 0-5. Before being expressed in MGCP, the MOS score is multiplied by 10 and any fractional part is truncated. This results in an integer with valid values in the range 10 -50.
Packet loss concealment type	PLC	The type of packet loss concealment algorithm in use; should be one of the following enumeration values: 0 - Unspecified 1 - Disabled 2 - Enhanced 3 - Standard
Jitter Buffer Adaptive	JBA	The adaptability of the jitter buffer. Should be one of the following enumeration values: 0 - Unknown 1 - Reserved 2 - Non-adaptive 3 - Adaptive
Jitter Buffer Rate	JBR	The jitter buffer adjustment rate. Value is between 0 and 15.
Nominal jitter buffer delay	JBN	Current nominal delay in milliseconds that corresponds to the nominal jitter buffer delay for packets that arrive exactly on time.
Maximum jitter buffer delay	JBM	Current maximum delay in milliseconds that corresponds to the earliest arriving packet that would not be discarded. In simple queue implementations, this may correspond to the nominal jitter buffer delay. In adaptive jitter buffer implementations, this value may dynamically vary up to the absolute maximum jitter buffer delay (see below).
Absolute maximum jitter buffer delay	JBS	Absolute maximum delay in milliseconds that an adaptive jitter buffer can reach under worst case conditions. For fixed jitter buffers, this must be set to the maximum jitter buffer delay.
Inter-arrival Jitter	IAJ	Using definition from RFC 3550, interarrival jitter is defined to be the mean deviation (smoothed absolute value) of the difference in packet spacing at the receiver compared to the sender for a pair of packets. This metric is provided in milliseconds

2.3.2 Extension Parameter Grammar

The following describes the ABNF encoding for the local and remote metrics parameters.

```
LocalVoIPMetrics = "XRM/LVM:" 0*WSP [VoIPMetric 0*("," 0*WSP VoIPMetric)]
```

```
RemoteVoIPMetrics = "XRM/RVM:" 0*WSP [VoIPMetric 0*("," 0*WSP VoIPMetric)]
```

```

VoIPMetric      = NetworkPacketLossRate
                  / JitterBufferDiscardRate
                  / BurstLossDensity
                  / GapLossDensity
                  / BurstDuration
                  / GapDuration
                  / RoundTripNetworkDelay
                  / EndSystemDelay
                  / SignalLevel
                  / NoiseLevel
                  / ResidualEchoReturnLoss
                  / MinimumGapThreshold
                  / RFactor
                  / ExternalRFactor
                  / EstimatedMOS-LQ
                  / EstimatedMOS-CQ
/ PacketLossConcealmentType
                  / JitterBufferAdaptive
                  / JitterBufferRate
                  / JitterBufferNominal
                  / JitterBufferMax
                  / JitterBufferAbsMax
                  / InterArrivalJitter

NetworkPacketLossRate      = "NLR=" 1*3(DIGIT) ;0-255

JitterBufferDiscardRate    = "JDR=" 1*3(DIGIT) ;0-255

BurstLossDensity           = "BLD=" 1*3(DIGIT) ; 0-255

GapLossDensity             = "GLD=" 1*3(DIGIT) ; 0-255

BurstDuration              = "BD=" 1*5(DIGIT) ; 0-65535

GapDuration                = "GD=" 1*5(DIGIT) ; 0-65535

MinimumGapThreshold        = "GMN=" 1*3(DIGIT) ; 1-255

RoundTripNetworkDelay      = "RTD=" 1*5(DIGIT) ;0-65535

EndSystemDelay             = "ESD=" 1*5(DIGIT) ;0-65535

SignalLevel                = "SL=" ["-"] 1*3(DIGIT) ; -128 to 127

NoiseLevel                 = "NL=" ["-"] 1*3(DIGIT) ; -128 to 127

ResidualEchoReturnLoss     = "RERL=" 1*3(DIGIT) ;0-127

RFactor                    = "RF=" 1*3(DIGIT) ;0-100, or 127

ExternalRFactor            = "XRF=" 1*3(DIGIT) ;0-100, or 127

EstimatedMOS-LQ            = "MLQ=" 1*3(DIGIT) ; 10-50, or 127

```

```

EstimatedMOS-CQ                = "MCQ=" 1*3(DIGIT) ; 10-50, or 127

PacketLossConcealmentType      = "PLC=" ("0" / "1" / "2" / "3")

JitterBufferAdaptive           = "JBA=" ("0" / "1" / "2" / "3")

JitterBufferRate               = "JBR=" 1*2(DIGIT) ;0-15

JitterBufferNominal            = "JBN=" 1*5(DIGIT) ;0-65535

JitterBufferMax                = "JBM=" 1*5(DIGIT) ;0-65535

JitterBufferAbsMax             = "JBS=" 1*5(DIGIT) ;0-65535

InterArrivalJitter             = "IAJ=" 1*5(DIGIT) ;0-65535

```

2.4 VoIP Metrics Negotiation

RFC 3611 [28] adds a new SDP attribute called "rtcp-xr" that can be used by a local endpoint of a connection to ask the remote endpoint to send RTCP XR reports for its locally collected metrics data. This attribute is encoded as follows:

```
rtcp-xr:<xr-format>
```

This rtcp-xr attribute can be used as either a session-level or media-level attribute. Only the voip-metrics xr-format option is mandated in PacketCable. Other options MAY be provided. When multiple options are provided they MUST be separated by a SPACE. See RFC 3611 [28] for more information on the use of this attribute parameter.

Send: The rtcp-xr:voip-metrics parameter MUST be sent in accordance with RFC 3611 [28] when voip-metrics is configured to be "on" via the Local Connection Options. When present, the parameter MAY appear as either a session-level or media-level attribute. This attribute parameter MUST NOT be sent when VoIP metrics is configured to be "off" or "negotiate". The endpoint MAY send other xr-format options independent of the VoIP Metrics LCO which are defined in RFC 3611 [28].

Receive: When the voip-metrics xr-format is received via this attribute parameter, the endpoint MUST include extended voip-metrics reports to the far end via RTCP as defined in RFC 3611 [28] unless the Call Agent has explicitly instructed the endpoint not to do so by configuring the VoIP metrics Local Connection Option to be "off". If the attribute parameter is not received, received but missing the voip-metrics option, or received but blank, i.e., no xr-format provided, then the endpoint MUST NOT include extended VoIP metrics reports to the far end via RTCP. Other xr-format options defined in RFC3611 [28] that are received MAY be ignored.

2.5 Implementation Considerations

As explained in the introduction, an endpoint or gateway must draw the distinction between collecting RTCP-XR voice quality metrics, responding to peer requests for metrics and reporting the metrics to the local call agent. The collection of metrics can be triggered by one or both of the following:

- A remote request from a peer gateway (via SDP)
- A local request from the controlling call agent (via a

CreateConnection or ModifyConnection)

The gateway MUST report the metrics during a DeleteConnection or AuditConnection (where VoIP metrics parameters are audited) sequence if the most recent CreateConnection or ModifyConnection enabled metric reporting (xrm/mcr: on).

If a VoIP metric is not computed locally, it MUST be omitted in the XRM/LVM: line. However, an endpoint supporting the XRM package MUST compute and report at least one VoIP metric from the table above.

If the remote end indicates that a particular RTCP-XR parameter is unavailable by sending a dummy value, then the dummy value MUST be reported or the parameter MUST be omitted from the XRM/RVM: line. If the remote end does not provide any remote VoIP metrics data, then the gateway response to a request for metrics data is based on the triggering command as follows:

- DeleteConnection: the XRM/RVM: parameter MUST be omitted,
- AuditConnection: the XRM/RVM: parameter MUST be included as an empty parameter.

3. Call Flow Examples

The following examples are partial call flows for the local part of a call that is invoking the RTCP-XR feature.

3.1 Call agent initiated DLCX

Step 1:

The Call Agent issues a CreateConnection command to the gateway instructing it to use PCMU media encoding and to invoke RTCP-XR metric collection, responding and reporting:

```
CRCX 1000 ds/dsl-1/1@gw-o.example.net MGCP 1.0
C: 1
L: a:PCMU, xrm/mcr: on
M: recvonly
```

Step 2:

The gateway acknowledges the command and includes the rtxp-xr attribute set to "voip-metrics":

```
200 1000 OK
I:1

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
a=rtcp-xr:voip-metrics
```

Step 3:

The gateway receives a ModifyConnection command containing an LCO with no xrm/mcr parameter, which indicates that the endpoint should continue to use the previously received value of "on". Also, the received SDP contains no rtcp-xr parameter, which indicates that the remote endpoint does not want to receive RTCP XR reports from this endpoint. Note that the absence of the rtcp-xr attribute in the received SDP does not signal the remote endpoint's willingness or ability to honor a previous rtcp-xr"voip-metrics" request it received from this endpoint.

```
MDCX 1001 ds/ds1-1/1@gw-o.example.net MGCP 1.0
C: 1
L: a:PCMU,G728
M: recvonly
```

```
v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
s=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 3456 RTP/AVP 8
```

Step 4:

The gateway acknowledges the command and, since VoIP metrics collection/reporting is still "on", includes the rtcp-xr attribute set to "voip-metrics":

```
200 1001 OK
I:1

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 8
a=rtcp-xr:voip-metrics
```

Step 5:

The Call Agent eventually issues a DeleteConnection:

```
DLCX 1100 ds/ds1-1/1@gw-o.example.net MGCP 1.0
C: 1
I: 1
```

Step 6:

The gateway acknowledges the command and returns the standard MGCP connection parameters along with the RTCP-XR parameters (note that new-lines shown here are for document formatting reasons only):

```
250 1100 OK
P: PS=1000, OS=60000, PR=800, OR=50000, PL=10, JI=27, LA=48,
```

```

XRM/LVM: NLR=128, JDR=40, BLD=33, GLD=10, BD=55, GD=45, RTD=1000,
ESD=1000, SL=-10, NL=2, RERL=23, GMN=80, RF=19, XRF=19, MLQ=19,
MCQ=22, PLC=1, JBA=2, JBR=8, JBN=40, JBM=80, JBS=120, IAJ=10
XRM/RVM: NLR=128, JDR=40, BLD=33, GLD=10, BD=55, GD=45, RTD=1000,
ESD=1000, SL=-10, NL=2, RERL=23, GMN=80, RF=19, XRF=19, MLQ=19,
MCQ=22, PLC=1, JBA=0, JBR=8, JBN=30, JBM=60, JBS=100, IAJ=15

```

3.2 Gateway initiated DLCX

Step 1:

The Call Agent issues a CreateConnection command to the gateway instructing it to use PCMU media encoding and to invoke RTCP-XR metric collection and reporting:

```

CRCX 1000 ds/dsl-1/1@gw-o.example.net MGCP 1.0
C: 1
L: a:PCMU, xrm/mcr: on
M: recvonly

```

Step 2:

The gateway acknowledges the command and includes the rtxp-xr attribute set to "voip-metrics":

```

200 1000 OK
I:1

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
a=rtcp-xr:voip-metrics

```

Step 3:

The Gateway issues a DeleteConnection command including the connection parameters and the RTCP metrics:

```

DLCX 1100 ds/dsl-1/1@gw-o.example.net MGCP 1.0
C: 1
I: 1
E: 900 - Hardware error
    P: PS=1000, OS=60000, PR=800, OR=50000, PL=10, JI=27, LA=48,
XRM/LVM: NLR=128, JDR=40, BLD=33, GLD=10, BD=55, GD=45, RTD=1000,
ESD=1000, SL=-10, NL=2, RERL=23, GMN=80, RF=19, XRF=19, MLQ=19,
MCQ=22, PLC=1, JBA=2, JBR=8, JBN=40, JBM=80, JBS=120
XRM/RVM: NLR=128, JDR=40, BLD=33, GLD=10, BD=55, GD=45, RTD=1000,
ESD=1000, SL=-10, NL=2, RERL=23, GMN=80, RF=19, XRF=19, MLQ=19,
MCQ=22, PLC=1, JBA=0, JBR=8, JBN=30, JBM=60, JBS=100

```

Note that the IAJ parameter is not included in the XRM/LVM or the XRM/RVM lines. This is just an example of allowed parameter omission. This example could have very well been constructed with the IAJ parameter.

Step 4:

The Call Agent acknowledges the command:

```
200 1100 OK
```

4. Security Considerations

This package inherits the security considerations of the base MGCP protocol. A possible new security issue involves a remote endpoint placing extra burden on a local endpoint by requesting metric collection on that local endpoint. This threat can be mitigated by the call agent specifically prohibiting the generation of local metrics using the local connection option: xrm/mcr: off.

5. IANA Considerations

The IANA is hereby requested to register the following MGCP package:

Package Title	Name	Version
-----	----	-----
RTCP-XR metrics	XRM	0

6. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3435] F. Andreassen, B. Foster, "Media Gateway Control Protocol (MGCP) Version 1.0", RFC 3435, January 2003.

[RFC3611] T. Friedman, R. Caceres, A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.

7. Acknowledgments

Contributors to this document include Flemming Andreassen, David Ensign, and the PacketCable NCS Specification Team.

A.5 Vbd and GPMD packages

FM package as specified in RFC 3660 MUST be implemented for use in conjunction with V.152 [30] as defined in the VBD package below.

Media Gateway Control Protocol Voiceband Data Package and
General Purpose Media Descriptor Parameter Package

Abstract

This document defines Media Gateway Control Protocol (MGCP) packages that enable a Call Agent to authorize and monitor the transition of a connection to and from voiceband data (VBD) with or without redundancy and FEC (forward error correction). Although the focus is on VBD, the General-Purpose Media Descriptor Parameter package can be used to authorize other modes of operation, not relevant to VBD, for a particular codec. In addition to the definition of these new packages, this document describes the use of the Media Format Parameter package and Fax package with VBD, redundancy and FEC.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC-2119 [RFC2119].

Table of Contents

1. Introduction.....	1
2. Voiceband Data Package Definition.....	3
2.1 Events and Signals.....	3
2.1.1 Gateway Controlled Voiceband Data.....	3
2.1.2 No Negotiated Procedure for Voiceband Data.....	8
3. General-Purpose Media Descriptor Parameter Package Definition.....	10
3.1 LocalConnectionOptions.....	10
3.1.1 General-Purpose Media Descriptor Parameter.....	11
4. Use of Media Format Parameter Package with VBD and Redundancy.....	13
5. Use of Media Format Parameter Package with VBD and FEC.....	14
6. Use of Fax Package with VBD.....	15
7. Call Flow Examples.....	18
7.1 Modem Call with Gateway Controlled VBD.....	18
7.2 Fax Call with Gateway Controlled VBD and Call Agent Controlled T.38.....	22
8. Security Considerations.....	29
9. IANA Considerations.....	29
10. Summary of Changes.....	30
11. Normative References.....	30
12. Informative References.....	31
13. Acknowledgements.....	31
14. Author's Address.....	31

1.

Introduction

The term voiceband data (or simply VBD) refers to the use of a suitable

voiceband codec (commonly G.711u or G.711a) for the transport of data payloads using the RTP protocol defined in IETF RFC 3550 [RFC3550]. This document defines Media Gateway Control Protocol (MGCP) [RFC3435] packages that enable a Call Agent to authorize and monitor the transition of a connection to and from VBD with or without redundancy [RFC2198] and FEC (forward error correction) [RFC2733].

There are a number of different VBD procedures. These procedures vary in terms of how the transition to and from VBD is coordinated end to end. Some coordination techniques are mutually negotiated by the two Gateways using the SDP [RFC2327]. These coordination techniques include,

- ITU-T Recommendation V.150.1 State Signalling Event (SSE) [V1501]
- ITU-T Recommendation V.152 Payload Type Switching [V152]

Other coordination techniques are not negotiated. For example, the detection of fax, modem and text tones in the direction from the IP to the GSTN network MAY result in switch to VBD or a change (e.g., disable echo cancellation) to the Gateway controlled VBD procedure already in place. The IP-side detected tone serves as both a VBD stimulus and a coordination technique.

IETF RFC 2833 [RFC2833] can be used to convey fax and modem events and tones. As with IP-side tone detection, the telephone event MAY serve as both a VBD stimulus and a coordination technique. Note that while the use of RFC 2833 to convey fax and modem events and tones is negotiated, the use of RFC 2833 as a Gateway VBD coordination technique (at present) is not.

The Voiceband Data (VBD) package is defined to support all VBD procedures. This document does not address the relative merits of different procedures nor advocate one procedure over another.

We will use the term VBD to refer to Voiceband Data in general. In referring to VBD the package, we will use the term VBD package. We use the term "audio" (with double quotes) to refer to the IANA media type. We use the term audio (without double quotes) to refer to the use of the "audio" media type for (most commonly) voice.

A package is defined for the General-Purpose Media Descriptor Parameter [V152]. In the context of VBD, the General-Purpose Media Descriptor Parameter (GPMD) package is used to authorize the negotiation of a particular codec for use with VBD. The General-Purpose Media Descriptor Parameter is "general" in nature and may be used in applications other than VBD.

The Media Format Parameter (FM) package [FMPKG] describes the use of the standard audio MIME subtype "RED" in conjunction with the "fmt" LocalConnectionOption in order to authorize the negotiation of redundancy [RFC2198], to identify the levels of redundancy and the media format associated with each redundancy level. This document will further explore the use of the FM package with VBD and redundancy.

This document will explore the use of the FM package with VBD and FEC.

The VBD package is intended to complement the MGCP Fax (FXR) package [FAXPKG]. This document will explore the use of the FXR package with VBD.

The VBD package definition is provided in section 2. The GPMD package definition is provided in section 3. In section 4, we discuss the use of the FM package with VBD and redundancy. In section 5, we discuss the use of the FM package with VBD and FEC. In section 6, we discuss the use of the FXR package with VBD. In section 7, we provide two call flow examples showing how to use the VBD and GPMD packages. Security considerations are found in section 8, followed by the IANA considerations and references.

2.

Voiceband Data Package Definition

A package is defined for (Voiceband Data) VBD. The package defines new events as detailed below.

Package Name: VBD
Package Version: 0

2.1

Events and Signals

The following events are defined in support of the above:

Symbol	Definition	R	S	Duration
gwvbd	Gateway Controlled VBD	x		
nopvbd	No Negotiated Procedure for VBD	x		

The definitions of the individual events are provided in the following subsections.

2.1.1 Gateway Controlled Voiceband Data

The "Gateway Controlled Voiceband Data" (or simply "gwvbd") event occurs when a gwvbd procedure has been negotiated and VBD stimulus is detected. The "gwvbd" event may occur when the gwvbd procedure is updated (e.g., upon detecting new stimulus) and when the procedure fails. The "gwvbd" event occurs when the gwvbd procedure ends. The gwvbd procedure MUST be negotiated with the other side by passing and recognizing relevant parameters via the LocalConnectionDescriptor and RemoteConnectionDescriptor.

The following recommendations from MGCP [RFC3435] apply,

In this section, we provide a formal description of the protocol syntax, following the "Augmented BNF for Syntax Specifications" defined in RFC 2234. The syntax makes use of the core rules defined in RFC 2234, Section 6.1, which are not included here. Furthermore, the syntax follows the case-sensitivity rules of RFC 2234, i.e., MGCP is case-insensitive (but SDP is not). It should be noted, that ABNF does not provide for implicit specification of linear white space and MGCP messages MUST thus follow the explicit linear white space rules provided in the grammar below. However, in line with general

robustness principles, implementers are strongly encouraged to tolerate additional linear white space in messages received.

The RequestedEvent is encoded as,

```
GwVbdReqEvent = "gwvbd"
```

The ObservedEvent is encoded as,

```
GwVbdObsEvent = GwVbdObsEventStart / GwVbdObsEventUpdate /
                  GwVbdObsEventStop / GwVbdObsEventFailure

GwVbdObsEventStart  = "gwvbd(start" Rc [Codec] [Coord] [Dir] ")"
GwVbdObsEventUpdate = "gwvbd(update" Rc [Codec] [Dir] ")"
GwVbdObsEventStop   = "gwvbd(stop" [Rc] [Codec] ")"
GwVbdObsEventFailure = "gwvbd(failure" [Rc] [Codec] ")"

Codec = "," *WSP "codec=" CodecString
CodecString = (ALPHA / DIGIT) *(ALPHA / DIGIT / "-" / "_" / "." / "/")
Coord = "," *WSP "coord=" CoordinationTechnique
CoordinationTechnique = "v152ptsw" / "v150fw"
Rc = "," *WSP "rc=" ReasonCode
ReasonCode = 1*(ALPHA / DIGIT / "-" / "_" / "." / "/")
              ; Refer to the values listed in the tables below.
Dir = "," *WSP "dir=" Direction
Direction = "GstnToIp" / "IpToGstn"
```

ABNF does not provide for position independent parameters. The "rc", "codec", "coord", and "dir" parameters, if present, MUST appear in the relative order shown.

The "start", "update", "stop" and "failure" ObservedEvent parameters are defined:

* VBD Start (start)

The gwvbd procedure was initiated. The Call Agent SHOULD refrain from issuing media handling instructions to the Gateway until either a "gwvbd(stop)" or "gwvbd(failure)" event is generated. One and only one "gwvbd(stop)" or "gwvbd(failure)" event is generated corresponding to each "gwvbd(start)" event.

* VBD Update (update)

The gwvbd procedure was updated. The "gwvbd(update)" event MUST only be generated after a "gwvbd(start)" event and before a "gwvbd(stop)" or "gwvbd(failure)" event.

* VBD Stop (stop)

The gwvbd procedure ended and the Gateway did not detect any errors. Note that this does not necessarily imply a successful fax, modem, or text transmission. It merely indicates that the gwvbd procedure has ended and the procedure itself did not encounter any errors. The "stop" parameter may

correspond to a change from VBD to a non-VBD "audio" codec or from VBD to another media type such as "image" or "text". This change may be under Call Agent or Gateway control. For example, the Gateway may coordinate the switch from VBD to "image/t38" through the exchange of SSEs [T38 and V152]. For an example involving Call Agent control, refer to the "MC" Reason Code. In both examples, the gwvbd procedure ends with the media change.

* VBD Failure (failure)

The gwvbd procedure ended abnormally. Some kind of problem was encountered in the gwvbd procedure and the procedure ended.

One of the above parameters will be present when the event is reported. The "gwvbd" event MAY be parameterized with additional parameters in ObservedEvents (including the parameters defined below), however exactly one of the "start", "update", "stop", or "failure" parameters MUST be present and MUST be the first parameter supplied.

Call Agents and Gateways MUST implement the "start" and "stop" parameters and MAY implement the "update" and "failure" parameters. Call Agents MAY and Gateways MUST implement the "rc" parameter in conjunction with the "start" and "update" parameters. A Call Agent MUST ignore all unknown ObservedEvent parameters including parameters which are defined as part of this specification and not implemented.

The "rc", "codec", "coord" and "dir" ObservedEvent parameters are defined:

* Coordination Technique (coord=<CoordinationTechnique>)

The technique used to coordinate the transition to and from VBD with the remote endpoint. The Coordination Techniques are summarized in the following table:

CoordinationTechnique	Description
v152ptsw	V.152 Payload Type Switching
v150fw	V.150.1 SSE

With the "v152ptsw" Coordination Technique, payload type switching [V152] is used to coordinate the transition to and from VBD.

With the "v150fw" Coordination Technique, state signalling events [V1501] are used to coordinate the transition to and from VBD.

The list of Coordination Techniques may be extended to include values with meaning mutually understood between the Gateway and the Call Agent.

* Reason Code (rc=<ReasonCode>)

With the "start" and "update" parameter, the reason for triggering the switch/change to VBD. With the "stop" and "failure" parameter, the reason for triggering the switch from VBD. The Reason Codes in the following

table, which are based on the ITU-T FAX/Textphone/Modem Tones Detection package [H2482], ITU-T V.150.1 Amendment 1 [V1501A1] and ITU-T V.152 [V152], may be used with the "start" and "update" parameter:

ReasonCode	Description
CNG	T.30 fax calling
V21flag	V.21 tone and flags for fax answering
CIV18	V.8 CI with V.18 call function
XCI	V.18 XCI
V18txp	V.18 txp
Belltone	Bell 103 carrier, high or low frequency channel (ITU-T Recommendation V.18)
Baudot	Baudot initial tone and character (ITU-T Recommendation V.18)
Edt	EDT initial tone and character (ITU-T Recommendation V.18)
CIdata	V.8 CI with any data call function
CT	V.25 calling tone
CIfax	V.8 CI with facsimile call function
V21tone	V.21 carrier, high or low frequency channel
V23tone	V.23 carrier, high or low frequency channel
V8bis	V.8 bis modem handshaking signal
ANS	V.25 ANS, equivalent to T.30 CED from answering terminal
/ANS	V.25 ANS with periodic phase reversals
ANSam	V.8 ANSam
/ANSam	V.8 ANSam with periodic phase reversals
CMFax	V.8 CM sequence indicating fax call function
JMFax	V.8 JM sequence indicating fax call function
CMData	V.8 CM sequence indicating unspecified data call function
JMData	V.8 JM sequence indicating unspecified data call function
CMText	V.8 CM sequence indicating text call function
JMText	V.8 JM sequence indicating text call function
PTSW	Payload type switch as defined in V.152

For solutions involving textphones using a modulation with interspersed Text and speech on the same "channel" such as Baudot and EDT, the Call Agent SHOULD interpret the ReasonCode parameter as part of the "vbd/gwvbd(start)" event in order to differentiate between fax, modem and text. In the case of interspersed text and speech the Call Agent SHOULD remove the notification request for "vbd/gwvbd" upon receiving the "vbd/gwvbd(start)" event in order to avoid large numbers of notifications. For example,

```
vbd/gwvbd(start, rc=Baudot)
```

With a ReasonCode of "PTSW", the Call Agent cannot differentiate text from fax/modem. In this case, the Call Agent SHOULD adopt a policy which guards against large numbers of notifications. We consider several such policies.

The Call Agent MAY remove the notification request for "vbd/gwvbd" upon receiving the "vbd/gwvbd(start, rc=PTSW)" event. With this policy, "update", "stop" and "failure" notifications will not be generated with text AND fax/modem.

The Call Agent MAY wait for a subsequent "vbd/gwvbd(update)" event which differentiates text from fax/modem. If the ReasonCode indicates interspersed text and speech, the Call Agent SHOULD remove the notification request for "vbd/gwvbd". For example,

```
vbd/gwvbd(update, rc=Edt)
```

The Call Agent MAY remove the notification request for "vbd/gwvbd" upon receiving a "vbd/gwvbd(stop)" event without having differentiated between text and fax/modem.

The Call Agent MAY remove the notification request for "vbd/gwvbd" after having received a number of "vbd/gwvbd(start)" events without having differentiated between text and fax/modem. The specific number of events after which the notification request is removed is considered an implementation detail outside the scope of this specification.

Reason Codes applicable with the "stop" parameter:

ReasonCode	Description
SIL	Bidirectional silence
Voice	Voice signals
PTSW	Payload type switch as defined in V.152
MC	Media change

The "MC" Reason Code indicates that the media type has changed from "audio" (to "image", "text", ...) or the "audio" media format has changed from a VBD codec (for a reason other than "PTSW"). For example, the gwvbd procedure may be initiated upon detecting CED. Subsequently, the Call Agent controlled T.38 procedure of the MGCP Fax (FXR) package [FAXPKG] may be initiated upon detecting V.21 flags. Upon receipt of a "t38(start)" event, the Call Agent will instruct the Gateway to switch from VBD to T.38 through the use of a ModifyConnection command involving a LocalConnectionOption encoding method of "L:a:image/t38" and/or a RemoteConnectionDescriptor with an "image/t38" media description. This stops the gwvbd procedure. There is no specific interdependency between the VBD package and the FXR package (or any other package). The gwvbd procedure is stopped as a consequence of the media change, not as a direct consequence of the T.38 procedure being initiated. Note that in this situation the "t38(start)" event will be sent before the "gwvbd(stop)" event. The Call Agent MAY choose to infer that the gwvbd procedure has ended upon receiving the "t38(start)" event and disable the notification of the "gwvbd" event. Refer to the example call flow in section 7.2.

Reason Codes applicable with the "failure" parameter:

ReasonCode	Description

TO	Indicates that a timeout has occurred
----	---------------------------------------

The list of Reason Codes may be extended to include values with meaning mutually understood between the Gateway and the Call Agent.

* Codec String (codec=<CodecString>)

With the "start" and "update" parameter, the MIME type associated with the switch/change to VBD (e.g., "audio/RED", "audio/PCMU", "audio/PCMA", "audio/G726-32", "audio/clearmode", ...). With the "stop" and "failure" parameter, the MIME type associated with the switch from VBD (e.g., "audio/G729", "image/t38", "text/t140", "audio/v150mr", ...). These strings should be full MIME types as listed in,

<http://www.iana.org/assignments/media-types>

* Direction of Stimulus (dir=<Direction>)

With the "start" and "update" parameter, the direction of the stimulus which resulted in the switch/change to VBD.

Direction	Description
GstnToIp	Stimulus detected from GSTN to IP
IpToGstn	Stimulus detected from IP to GSTN

Call Agents and Gateways MAY implement the "coord", "rc", "codec" and "dir" parameters. A Call Agent MUST ignore all unknown ObservedEvent parameters including parameters which are defined as part of this specification and not implemented.

2.1.1.1 Gateway Controlled Voiceband Data Examples

The following examples illustrate the encoding of the "gwvbd(start)" event,

```
O: vbd/gwvbd(start, rc=ANS)
O: vbd/gwvbd(start, rc=ANS, codec=audio/PCMU, coord=v152ptsw)
O: vbd/gwvbd(start, rc=PTSW, codec=audio/RED)
```

The following example illustrates the encoding of the "gwvbd(update)" event,

```
O: vbd/gwvbd(update, rc=/ANSam, dir=IpToGstn)
```

The following examples illustrate the encoding of the "gwvbd(stop)" event,

```
O: vbd/gwvbd(stop)
O: vbd/gwvbd(stop, rc=SIL, codec=audio/G729)
O: vbd/gwvbd(stop, rc=MC, codec=image/t38)
```

The following examples illustrate the encoding of the "gwvbd(failure)" event,

```
O: vbd/gwvbd(failure, codec=audio/G729)
O: vbd/gwvbd(failure, rc=TO, codec=audio/G729)
```

2.1.1.2 No Negotiated Procedure for Voiceband Data

The "No Negotiated Procedure for Voiceband Data" (or simply "nopvbd") event occurs when a VBD procedure has not been negotiated and VBD stimulus is detected. The "nopvbd" event may occur when the procedure is updated (e.g., upon detecting new stimulus), when the procedure ends and when the procedure fails.

The nopvbd procedure MAY involve VBD handling including but not limited to adjusting gain and jitter, disabling voice activity detection and DC offset filters. The nopvbd procedure MAY involve switching to another codec. The Call Agent MAY have to issue further commands in response to the "nopvbd" event in order to ensure a successful VBD call.

As with the "gwvbd" event, the same recommendations from MGCP [RFC3435] regarding ABNF, general robustness principles and white space apply.

The RequestedEvent is encoded as,

```
NopVbdReqEvent = "nopvbd"
```

The ObservedEvent is encoded as,

```
NopVbdObsEvent = NopVbdObsEventStart / NopVbdObsEventUpdate /
                  NopVbdObsEventStop / NopVbdObsEventFailure

NopVbdObsEventStart   = "nopvbd(start" Rc [Codec] [Dir] ")"
NopVbdObsEventUpdate  = "nopvbd(update" Rc [Codec] [Dir] ")"
NopVbdObsEventStop    = "nopvbd(stop" [Rc] [Codec] ")"
NopVbdObsEventFailure = "nopvbd(failure" [Rc] [Codec] ")"

; The following ABNF notation is common with the "gwvbd" ObservedEvent.
Codec = "," *WSP "codec=" CodecString
CodecString = (ALPHA / DIGIT) *(ALPHA / DIGIT / "-" / "_" / "." / "/" )
Rc = "," *WSP "rc=" ReasonCode
ReasonCode = 1*(ALPHA / DIGIT / "-" / "_" / "." / "/" )
              ; Refer to the values listed in the tables below.
Dir = "," *WSP "dir=" Direction
Direction = "GstnToIp" / "IpToGstn"
```

ABNF does not provide for position independent parameters. The "rc", "codec", and "dir" parameters, if present, MUST appear in the relative order shown.

The "start", "update", "stop" and "failure" ObservedEvent parameters are defined:

* Start VBD (start)

The nopvbd procedure was initiated. The Call Agent may have to issue further commands in order to ensure a successful VBD call (e.g., switch to another codec). At most one "nopvbd(stop)" or "nopvbd(failure)" event MAY

be generated corresponding to each "nopvbd(start)" event. The Call Agent MAY need to infer that the nopvbd procedure has ended.

* VBD Update (update)

The nopvbd procedure was updated. The "nopvbd(update)" event MUST only be generated after a "nopvbd(start)" event and before a "nopvbd(stop)" or "nopvbd(failure)" event.

* VBD Stop (stop)

The nopvbd procedure ended and the Gateway did not detect any errors. Note that this does not necessarily imply a successful fax, modem, or text transmission. It merely indicates that the nopvbd procedure has ended and the procedure itself did not encounter any errors. Refer to the definition of the "stop" parameter from the "gwvbd" event in section 2.1.1 for additional information.

* VBD Failure (failure)

The nopvbd procedure ended abnormally. Some kind of problem was encountered in the nopvbd procedure and the procedure ended.

Call Agents and Gateways MUST implement the "start" parameter and MAY implement the "update", "stop" and "failure" parameters. Call Agents MAY and Gateways MUST implement the "rc" parameter in conjunction with the "start" and "update" parameters. A Call Agent MUST ignore all unknown ObservedEvent parameters including parameters which are defined as part of this specification and not implemented.

The definitions of the "rc", "codec" and "dir" ObservedEvent parameters are taken from the "gwvbd" event.

As with the "gwvbd" event, the same recommendations regarding interspersed text and speech apply.

2.1.2.1 No Negotiated Procedure for Voiceband Data Examples

The following examples illustrate the encoding of the "nopvbd(start)" event,

```
O: vbd/nopvbd(start, rc=ANS)
O: vbd/nopvbd(start, rc=ANS, codec=audio/PCMU)
```

The following example illustrates the encoding of the "nopvbd(update)" event,

```
O: vbd/nopvbd(update, rc=/ANSam, dir=IpToGstn)
```

The following examples illustrate the encoding of the "nopvbd(stop)" event,

```
O: vbd/nopvbd(stop)
O: vbd/nopvbd(stop, rc=SIL, codec=audio/G729)
O: vbd/nopvbd(stop, rc=MC, codec=image/t38)
```

The following examples illustrate the encoding of the "nopvbd(failure)" event,

```
O: vbd/nopvbd(failure, codec=audio/G729)
O: vbd/nopvbd(failure, rc=TO, codec=audio/G729)
```

3.

General-Purpose Media Descriptor Parameter Package Definition

A package is defined for the General-Purpose Media Descriptor Parameter [V152]. The package defines a new LocalConnectionOption as detailed below.

```
Package Name:      GPMD
Package Version:   0
```

3.1

LocalConnectionOptions

The following LocalConnectionOption is defined in support of the above:

Symbol	Definition
-----	-----
gpmd	General-Purpose Media Descriptor Parameter

The definition of the LocalConnectionOption is provided in the following subsection.

3.1.1 General-Purpose Media Descriptor Parameter

The General-Purpose Media Descriptor Parameter LocalConnectionOption is similar to the "gpmd" SDP [RFC2327] attribute defined in ITU-T Recommendation V.152 [V152] and is applicable to all of the same media formats that the corresponding SDP "gpmd" attribute could be used with.

The General-Purpose Media Descriptor Parameter is encoded as the keyword "gpmd" or "o-gpmd", followed by a colon and a quoted string beginning with the media format name (MIME subtype only) followed by a space, followed by the media format parameters associated with that media format,

```
gpmd/gpmd:"<format> <parameter list>"
```

For simplicity, we will use the terms "codec" and "media format" interchangeably in the following. Multiple media formats may be indicated by either repeating the "gpmd" LocalConnectionOption multiple times, such as:

```
L: a:codecl;codec2, gpmd/gpmd:"codecl parameterX",
    gpmd/gpmd:"codec2 parameterY"
```

or alternatively by having a single "gpmd" keyword followed by a colon, and a semi-colon separated list of quoted strings for each General-Purpose Media Descriptor Parameter, as in:

```
L: a:codecl1;codecl2, gpmd/gpmd:"codecl1 parameterX";  
    "codecl2 parameterY"
```

The two formats may be mixed,

```
L: a:codecl1;codecl2;codecl3, gpmd/gpmd:"codecl1 parameterX",  
    gpmd/gpmd:"codecl2 parameterY";  
    "codecl3 parameterZ"
```

The carriage returns above are included for formatting reasons only and are not permissible in a real implementation. This holds true for all of the examples in this document.

If it is possible for the same codec to be requested with and without the "gpmd" parameter, the following could result:

```
L: a:codecl1;codecl1, gpmd/gpmd:"codecl1 parameterX"
```

However, it would not be clear if the "gpmd" parameter was to be applied to the first or the second occurrence of the codec. The problem with that is that codec ordering is important (i.e., codecs are listed in preferred order), and the above syntax does not provide a way to indicate if "parameterX" is preferred (i.e., associated with the first "codecl1") or not (i.e., associated with the second "codecl1"). In order to resolve this dilemma, the codec in the "gpmd" media format is followed by a colon and an <order>, where <order> is a number from one to N for occurrences of the same codec in the codec list i.e.:

```
L:a:codecl1;codecl1, gpmd/gpmd:"codecl1:2 parameterX"
```

indicates that "parameterX" is associated with the second instance of "codecl1" in the "a:codecl1;codecl1" list. If an invalid instance number is supplied (e.g., instance 3 where there are only two instances), then error code 524 - inconsistency in local connection options will be returned.

Pre-pending "gpmd" with the string "o-" (i.e., "o-gpmd") indicates that the parameter is optional. In that case, the Gateway may decide not to use the "gpmd" parameter specified, or only use it in part.

If the "gpmd" LocalConnectionOption parameter is not optional (i.e., does not have "o-" in front of it), and the LocalConnectionOption parameter value is either not recognized or not supported, then the associated codec is considered "not supported".

When auditing capabilities, the "gpmd" LocalConnectionOption parameter MUST be returned with a semi-colon separated list of supported formats and/or multiple independent "gpmd" parameters as in:

```
A: a:codecl1;codecl2, gpmd/gpmd:"codecl1 parameterX";  
    "codecl2 parameterY"
```

or,

```
A: a:codecl1;codecl1, gpmd/gpmd:"codecl1 parameterX"
```

One example uses the General-Purpose Media Descriptor Parameter LocalConnectionOption in conjunction with Gateway controlled Voiceband Data (or simply VBD) using payload type switching [V152]. In the context of VBD, the <format> must be an RTP/AVP payload type. The <parameter list> is a semicolon-separated list of "parameter=value" pairs,

```
L: a:codecl, gpmd/gpmd:"codecl parameterX=ValueA;parameterY=ValueB"
```

In the example below, G.729 is an audio codec and G.711u is a VBD codec,

```
L: a:G729;PCMU, gpmd/gpmd:"PCMU vbd=yes"
```

The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
m=audio 12345 RTP/AVP 18 96
a=rtpmap:96 PCMU/8000
a=gpmd:96 vbd=yes
```

If a request is made to audit the capabilities of an endpoint, and the endpoint supports G.711u as both an audio and VBD codec, then the "gpmd" LocalConnectionOption parameter might look like:

```
A: a:PCMU, p:10-40, e:on, s:on,
   m:sendonly;recvonly;sendrecv;inactive
A: a:PCMU, p:10-40, e:on, s:off,
   m:sendonly;recvonly;sendrecv;inactive,
   gpmd/gpmd:"PCMU vbd=yes"
```

Given that some parameters, such as e.g., silence suppression, are only compatible with G.711u as an audio codec, then the Gateway MUST return different capability sets corresponding to audio and VBD.

If we combine V.152 and redundancy [RFC2198], an example LocalConnectionOption might look as follows. In the example below, G.729 is an audio codec and G.711u is a VBD codec with a redundancy level of one:

```
L: a:G729;RED;PCMU, gpmd/gpmd:"PCMU vbd=yes", fntp:"RED PCMU/PCMU"
```

The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
m=audio 12345 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fnfnp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
```

Refer to section 4 for more examples involving V.152 and redundancy.

4.

Use of Media Format Parameter Package with VBD and Redundancy

The MGCP Media Format Parameter (FM) package [FMPKG] in conjunction with the standard audio MIME subtype "RED" may be used by the Call Agent to authorize the negotiation of redundancy [RFC2198], to identify the levels of redundancy and the media format associated with each redundancy level. An example of this was demonstrated in section 3.

The FM package states that the "fmt" LocalConnectionOption MUST be returned when auditing capabilities. Applying this to VBD and redundancy might result in:

```
A: a:PCMU, p:10-40, e:on, s:on,
    m:sendonly;recvonly;sendrecv;inactive
A: a:RED;PCMU, p:10-40, e:on, s:off,
    m:sendonly;recvonly;sendrecv;inactive,
    gpmd/gpmd:"PCMU vbd=yes",
    fmt:"RED PCMU/PCMU"
```

The FM package defines "instance syntax" in which,

```
L:a:codecl;codecl, fmt:"codecl:2 formatX"
```

indicates that "formatX" is associated with the second instance of "codecl" in the "a:codecl;codecl" list. The examples in the FM package are limited to the use of the instance syntax in conjunction with the media format. We propose the use of the instance syntax in conjunction with the media format parameters,

```
L:a:codecl;codec2;codec3;codec2, fmt:"codec3 codec2:2/codec2:2"
```

Let's build on the example of section 3. In the example below, G.729 is an audio codec, G.711u is both an audio codec and a VBD codec with a redundancy level of one:

```
L: a:G729;PCMU;RED;PCMU, gpmd/gpmd:"PCMU:2 vbd=yes",
    fmt:"RED PCMU:2/PCMU:2"
```

The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
m=audio 12345 RTP/AVP 18 0 96 97
a=rtpmap:96 RED/8000
a=fmt:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
```

Note that the relative preference of the LocalConnectionOption encoding methods is preserved in the "audio" media formats (i.e., payload types) as part of the media description. In this example, this reflects a preference for V.152 with redundancy versus without. No preference is inferred from the relative order of the different LocalConnectionOptions, namely "a", "gpmd/gpmd" and "fmt".

A Call Agent can authorize the negotiation of audio codecs and VBD codecs involving different levels of redundancy. In the example below, G.711u is a

VBD codec with a redundancy level of two (preferred) or one:

```
L: a:G729;RED;RED;PCMU, fmp:"RED PCMU/PCMU/PCMU", fmp:"RED:2 PCMU/PCMU",
    gpmd/gpmd:"PCMU vbd=yes"
```

The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
m=audio 12345 RTP/AVP 18 96 97 98
a=rtpmap:96 RED/8000
a=fmp:96 98/98/98
a=rtpmap:97 RED/8000
a=fmp:97 98/98
a=rtpmap:98 PCMU/8000
a=gpmd:98 vbd=yes
```

Redundancy can be applied to both audio codecs and VBD codecs. In the example below, G.729 is an audio codec with a redundancy level of two and G.711u is a VBD codec with a redundancy level of one:

```
L: a:RED;G729;RED;PCMU, fmp:"RED G729/G729/G729", fmp:"RED:2 PCMU/PCMU",
    gpmd/gpmd:"PCMU vbd=yes"
```

The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
m=audio 12345 RTP/AVP 96 18 97 98
a=rtpmap:96 RED/8000
a=fmp:96 18/18/18
a=rtpmap:97 RED/8000
a=fmp:97 98/98
a=rtpmap:98 PCMU/8000
a=gpmd:98 vbd=yes
```

5.

Use of Media Format Parameter Package with VBD and FEC

A Call Agent may authorize the negotiation of forward error correction (FEC) [RFC2733] with the standard audio MIME subtype "parityfec",

```
L: a:PCMU;parityfec
```

By default, we assume that FEC packets are to be sent as a separate stream. The corresponding media description in the SDP as part of the connection request acknowledgment might look like:

```
v=0
c=IN IP4 224.2.17.12
m=audio 49170 RTP/AVP 0 96
a=rtpmap:96 parityfec/8000
a=fmp:96 49172 IN IP4 224.2.17.12
```

If FEC is to be sent as a secondary codec in the redundant codec payload format [RFC2198], we again leverage the MGCP Media Format Parameter (FM)

package [FMPKG] in conjunction with the standard audio MIME subtype "RED",

```
L: a:G729;RED;PCMU;parityfec, gpmd/gpmd:"PCMU vbd=yes",
    fmt:"RED PCMU/parityfec"
```

The corresponding media description might look like:

```
v=0
c=IN IP4 224.2.17.12
m=audio 49170 RTP/AVP 18 96 97 98
a=rtpmap:96 RED/8000
a=fmt:96 97/98
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
a=rtpmap:98 parityfec/8000
```

The FM package states that the "fmt" LocalConnectionOption MUST be returned when auditing capabilities. Applying this to VBD, redundancy and FEC might result in:

```
A: a:PCMU, p:10-40, e:on, s:on,
    m:sendonly;recvonly;sendrecv;inactive
A: a:RED;PCMU;parityfec, p:10-40, e:on, s:off,
    m:sendonly;recvonly;sendrecv;inactive,
    gpmd/gpmd:"PCMU vbd=yes",
    fmt:"RED PCMU/parityfec"
```

6.

Use of Fax Package with VBD

The MGCP Fax (FXR) package [FAXPKG] is used by a Call Agent to authorize fax handling including Call Agent controlled T.38 and Gateway procedures such as V.152. With the FXR package, VBD falls into one of two categories: "special fax handling" as part of the Gateway procedure (resulting in the "gwfax" event); or "no special fax handling" as part of the Gateway and Off procedures (resulting in the "nopfax" event). In order for a VBD procedure to fall into the "special fax handling" category,

"... support for it MUST be negotiated with the other side by passing and recognizing relevant parameters via the LocalConnectionDescriptor and RemoteConnectionDescriptor."

A Gateway controlled VBD procedure such as V.152 MUST fall into the category of Gateway controlled mode involving "special fax handling". The resulting "gwfax" event is what informs the Call Agent to refrain from issuing media handling instructions which could otherwise have a negative impact on the Gateway procedure.

Consider the following example (with shorthand SDP notation):

```
CRCX 2000 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 1
M: sendrecv
L: a:G729;PCMU, gpmd/gpmd:"PCMU vbd=yes", fxr/fx:t38;gw
```

```
X: 1
R: fxr/t38, fxr/gwfax, fxr/nopfax
```

```
v=0
c=IN IP4 128.96.41.1
m=audio 3456 RTP/AVP 18 96
a=rtpmap:96 PCMU/8000
a=gpmd:96 vbd=yes
```

```
200 2000 OK
I: 1
```

```
v=0
m=audio 1296 RTP/AVP 18 96
a=rtpmap:96 PCMU/8000
a=gpmd:96 vbd=yes
```

The RemoteConnectionDescriptor does not indicate support for "image/t38" as a latent capability [RFC3407]. Consequently, the Gateway will not initiate the T.38 strict fax procedure, "t38", upon detecting fax stimulus (i.e., CNG, V.21 flags, ...). However, the two endpoints did successfully negotiate a Gateway controlled VBD procedure (e.g., V.152), therefore, Gateway controlled mode involving "special fax handling" is used. The "gwfax(start)" event will be generated upon detecting VBD (including fax) stimulus.

A Call Agent can express a preference for a Gateway procedure involving "special fax handling" over a T.38 procedure (strict or loose). For example,

```
L: fxr/fx:gw;t38
```

and,

```
L: fxr/fx:gw;t38-loose
```

However, with the existing syntax of the FXR package, a Call Agent cannot express a preference for one Gateway procedure over another, each with possibly different preferences relative to a T.38 procedure.

The FXR package allows a Gateway to implement additional fax handling parameters. We define just such a parameter by qualifying the existing "gw" parameter with a list of one or more MIME types,

```
Gateway = "gw[" mimeType 0*("|" mimeType) "]"
mimeType = mimeType "/" subType
; mimeType and subType from,
; http://www.iana.org/assignments/media-types/
```

By qualifying the "gw" parameter with a list of MIME types, we narrow the scope of the Gateway procedure. Consider the following examples in which the Call Agent authorizes the use of a Gateway controlled fax handling procedure:

- involving "image/t38" (e.g., T.38oUDPTL, T.38oTCP),

L: a:G729, fxr/fx:gw[image/t38]

- involving VBD (e.g., PCMU and V.152),

L: a:G729;PCMU, gpmd/gpmd:"PCMU vbd=yes", fxr/fx:gw[audio/PCMU]

- involving VBD with redundancy (e.g., PCMU, V.152 and RFC 2198),

L: a:G729;RED;PCMU, ftmp:"RED PCMU/PCMU", gpmd/gpmd:"PCMU vbd=yes",
fxr/fx:gw[audio/RED|audio/PCMU]

Only "special fax handling" involving one of the specified MIME types is authorized. Support for "special fax handling" involving one of the specified MIME types MUST be negotiated or this "instance" of the Gateway procedure is not initiated. Consider the following example in which the Call Agent authorizes the use of a Gateway controlled fax handling procedure:

- involving "audio/t38" (e.g., T.38oRTP),

L: a:G729;t38, fxr/fx:gw[audio/t38]

In this example, the call will fail if the Gateway fails to negotiate "audio/t38".

The "fx" LocalConnectionOption MAY now involve multiple instances of the "gw" parameter, each with a different list of MIME types. In order to authorize "no special fax handling", the Call Agent MUST include: the "gw" parameter without a MIME type; or the "off" parameter. The instance of the "gw" parameter without a MIME type should appear as the last instance of the "gw" parameter. In the following example,

L: a:G729;PCMU, fxr/fx:gw[image/t38];gw

the Call Agent authorizes the use of and expresses a preference for,

1. Gateway controlled image/t38 (e.g., T.38oUDPTL)
2. Any other Gateway procedure with "special fax handling"
3. No special fax handling

If present, the "off" parameter should appear as the last parameter. In the following example,

L: a:G729;PCMU;t38, fxr/fx:gw[audio/t38];off

the Call Agent authorizes the use of and expresses a preference for,

1. Gateway controlled audio/t38 (e.g., T.38oRTP)
2. No special fax handling

We can express relative preferences for different Gateway controlled fax handling procedures, not only with respect to one another, but with respect to T.38 procedures. Consider the following preferential list of fax handling procedures,

1. Gateway controlled audio/t38 (e.g., T.38oRTP)
2. Gateway controlled image/t38 (e.g., T.38oUDPTL)
3. Call Agent Controlled image/t38
4. Gateway controlled VBD with Redundancy (e.g., PCMU, V.152 and RFC 2198)
5. Gateway controlled VBD without Redundancy (e.g., PCMU and V.152)
6. Any other Gateway procedure with "special fax handling"
7. No special fax handling

This would be expressed as,

```
L: a:G729;PCMU;t38;RED;PCMU,
    gpmd/gpmd:"PCMU:2 vbd=yes",
    fntp:"RED PCMU:2/PCMU:2",
    fxr/fx:gw[audio/t38|image/t38];t38;gw[audio/RED|audio/PCMU:2];gw
```

Note that the bracketed form of the "gw" parameter is NOT defined as part of the VBD package. The bracketed form of the "gw" parameter is defined as an extension to the FXR package. Gateways that implement the bracketed form of the "gw" parameter MUST return this form of the parameter when capabilities are audited as illustrated by the following example:

```
A: fxr/fx:t38;t38-loose;gw[audio/t38|image/t38];gw:off
```

Support for the bracketed "gw" parameter MAY be spread across multiple capability lines,

```
A: a:RED;PCMU, p:10-40, e:on, s:off,
    m:sendonly;recvonly;sendrecv;inactive,
    gpmd/gpmd:"PCMU vbd=yes",
    fntp:"RED PCMU/PCMU",
    fxr/fx:gw[audio/RED|audio/PCMU]
A: a:t38, fxr/fx:gw[audio/t38]
A: a:image/t38, fxr/fx:t38;t38-loose;gw[image/t38]
```

A Call Agent SHOULD only attempt to leverage the bracketed form of the "gw" parameter in conjunction with an endpoint which indicates support for the bracketed syntax as part of the capabilities.

Call Agents and Gateways that do not support this form of the "gw" parameter MUST ignore the bracketed MIME type information consistent with the MGCP grammar [RFC3435].

7.

Call Flow Examples

In this section, we provide two example call flows. The first one illustrates a modem call under Gateway control using V.152. The second one illustrates a fax call under Gateway control using V.152 and Call Agent controlled T.38.

7.1

Modem Call with Gateway Controlled VBD

In this example, both sides support Gateway controlled VBD using V.152 with redundancy. We assume the originating and terminating Call Agent communicate

via the Session Initiation Protocol (SIP) [RFC3261]:

#	GW-o	CA-o	CA-t	GW-t
1		<- CRCX		
2	200(sdp-o) ->			
3		INVITE(sdp-o) ->		
4			CRCX(sdp-o) ->	
5				<- 200 (sdp-t)
6			<- 200(sdp-t)	
7	<- MDCX(sdp-t)			
8	200 ->			
9				<- ANS/T.30 CED
10				<- NTFY(gwvbd start)
11				200 ->
12	NTFY(gwvbd start) ->			
13	<- 200			
14				(modem ends)
15				<- NTFY(gwvbd stop)
16				200 ->
17	NTFY(gwvbd stop) ->			
18	<- 200			

Step 1:

The Call Agent issues a CreateConnection command to the Gateway instructing it to use G.729 media encoding and to notify it of the "gwvbd" and "nopvbd" events. The Call Agent authorizes the negotiation of G.711u as a VBD codec with a redundancy level of one:

```
CRCX 1000 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
L: a:G729;RED;PCMU, gpmd/gpmd:"PCMU vbd=yes", fntp:"RED PCMU/PCMU"
M: recvonly
R: vbd/gwvbd, vbd/nopvbd
X: 1
Q: process, loop
```

Step 2:

The Gateway acknowledges the command and includes SDP with codec information as well as V.152 and redundancy information:

```
200 1000 OK
I:1

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 18 96 97
a=rtpmap:96 RED/8000
a=fmtp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmid:97 vbd=yes

```

Step 3:

The originating Call Agent sends a SIP INVITE message with the SDP to the terminating Call Agent.

Step 4:

The terminating Call Agent issues a CreateConnection command to the terminating Gateway instructing it to use G.729 media encoding and to notify it of the "gwvbd" and "nopvbd" events. Again, the Call Agent authorizes the negotiation of G.711u as a VBD codec with a redundancy level of one:

```

CRCX 2000 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
L: a:G729;RED;PCMU, gpmid/gpmid:"PCMU vbd=yes", fmtp:"RED PCMU/PCMU"
M: sendrecv
R: vbd/gwvbd, vbd/nopvbd
X: 20
Q: process, loop

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 18 96 97
a=rtpmap:96 RED/8000
a=fmtp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmid:97 vbd=yes

```

Step 5:

The terminating Gateway supports V.152 and redundancy, and the RemoteConnectionDescriptor included indicates that the other side supports V.152 and redundancy. The terminating Gateway sends back a success response with its SDP which also includes V.152 and redundancy information:

```

200 2000 OK
I:2

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 18 96 97

```



```
a=rtpmap:96 RED/8000
a=fmtp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
```

Step 6:

The terminating Call Agent sends back a SIP 200 OK response to the originating Call Agent, which in turn sends a SIP ACK (not shown).

Step 7:

The originating Call Agent in turns sends a ModifyConnection command to the originating Gateway:

```
MDCX 1001 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
M: sendrecv

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=audio 1296 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fmtp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
```

Since the RemoteConnectionDescriptor indicates the other side supports V.152 and redundancy, the Gateway will in fact be able to use the Gateway controlled VBD procedure with redundancy. Had there not been any support for V.152 in the RemoteConnectionDescriptor, then this command would still have succeeded, however there would be no negotiated procedure for VBD handling.

Step 8:

The Gateway acknowledges the command. At this point, a call is established using G.729 encoding, and if a VBD call is detected, the Gateway controlled VBD procedure will be initiated.

Step 9-10:

A modem call now occurs. The terminating Gateway detects a T.30 CED tone (a.k.a. V.25 ANS) in the GSTN to IP direction and begins transmitting RTP packets with the negotiated redundant VBD payload type (96).

The "gwvbd(start)" event occurs and is notified to the Call Agent:

```
NTFY 2500 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: vbd/gwvbd(start, rc=ANS, codec=audio/RED, coord=v152ptsw)
X: 20
```

Step 11:

The Call Agent acknowledges the Notify command:

```
200 2500 OK
```

Step 12:

Upon receiving a RTP packet with the redundant VBD payload type (96), the originating Gateway begins transmitting RTP packets with the redundant VBD payload type.

The "gwvbd(start)" event occurs and is notified to the Call Agent:

```
NTFY 1500 ds/dsl-1/1@gw-0.whatever.net MGCP 1.0
O: vbd/gwvbd(start, rc=PTSW, codec=audio/RED)
X: 1
```

Step 13:

The Call Agent acknowledges the Notify command:

```
200 1500 OK
```

Step 14 - 15:

The modem call ends. The terminating Gateway detects bi-directional silence and begins transmitting RTP packets with the negotiated audio payload type (18).

The "gwvbd(stop)" event occurs and is notified to the Call Agent:

```
NTFY 2501 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: vbd/gwvbd(stop, rc=SIL, codec=audio/G729)
X: 20
```

Step 16:

The Call Agent acknowledges the Notify command:

```
200 2501 OK
```

Step 17:

Upon receiving a RTP packet with the audio payload type (18), the originating Gateway begins transmitting RTP packets with the audio payload type.

The "gwvbd(stop)" event occurs and is notified to the Call Agent:

```
NTFY 1501 ds/dsl-1/1@gw-0.whatever.net MGCP 1.0
O: vbd/gwvbd(stop, rc=PTSW, codec=audio/G729)
X: 1
```

Step 18:

The Call Agent acknowledges the Notify command:

200 1501 OK

The modem call is now over.

7.2

Fax Call with Gateway Controlled VBD and Call Agent Controlled T.38

In this example, both sides support Gateway controlled VBD using V.152 with Redundancy and Call Agent controlled T.38. We assume the originating and terminating Call Agent communicate via the Session Initiation Protocol (SIP) [RFC3261]:

#	GW-o	CA-o	CA-t	GW-t
1		<- CRCX		
2	200(sdp-o) ->			
3		INVITE(sdp-o) ->		
4			CRCX(sdp-o) ->	
5				<- 200 (sdp-t)
6			<- 200(sdp-t)	
7		<- MDCX(sdp-t)		
8	200 ->			
9				<- ANS/T.30 CED
10				<- NTFY(gwvbd start)
11				200 ->
12	NTFY(gwvbd start) ->			
13		<- 200		
14				<- V.21 Preamble
15				<- NTFY(t38 start)
16				200 ->
17				MDCX(t38) ->
18				<- 200(sdp-t2)
19			<- INVITE(sdp-t2)	
20		<- MDCX(sdp-t2)		
21	200(sdp-o2) ->			
22		200(sdp-o2) ->		
23			MDCX(sdp-o2) ->	
24				<- 200
25				(fax ends)
26				<- NTFY(t38 stop)
27				200 ->

Step 1:

The Call Agent issues a CreateConnection command to the Gateway instructing it to use G.729 media encoding and to use either the strict T.38 procedure or

the Gateway procedure. Consequently, the Call Agent requests notification of the "t38", "gwfax", "gwvbd" and "nopvbd" events. The Call Agent authorizes the negotiation of G.711u as a VBD codec with a redundancy level of one:

```
CRCX 1000 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
L: a:G729;RED;PCMU, gpmd/gpmd:"PCMU vbd=yes", fntp:"RED PCMU/PCMU",
  fxr/fx:t38;gw
M: recvonly
R: fxr/t38, fxr/gwfax, vbd/gwvbd, vbd/nopvbd
X: 1
Q: process, loop
```

Step 2:

The Gateway acknowledges the command and includes SDP with codec information as well as capability, V.152 and redundancy information:

```
200 1000 OK
I:1

v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
a=pmft: T38
m=audio 3456 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fntp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cdsc: 4 image udptl t38
```

Note that V.152 requires the use of the session-level "a=pmft" SDP attribute in order to express a preference for T.38 over V.152 for fax handling.

Step 3:

The originating Call Agent sends a SIP INVITE message with the SDP to the terminating Call Agent.

Step 4:

The terminating Call Agent issues a CreateConnection command to the terminating Gateway instructing it to use G.729 media encoding and to use either the strict T.38 procedure or the Gateway procedure. Consequently, the Call Agent requests notification of the "t38", "gwfax", "gwvbd" and "nopvbd" events. Again, the Call Agent authorizes the negotiation of G.711u as a VBD codec with a redundancy level of one:

```
CRCX 2000 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
```

```
C: 2
L: a:G729;RED;PCMU, gpmd/gpmd:"PCMU vbd=yes", fntp:"RED PCMU/PCMU",
  fxr/fx:t38;gw
M: sendrecv
R: fxr/t38, fxr/gwfax, vbd/gwvbd, vbd/nopvbd
X: 20
Q: process, loop

v=0
o=- 25678 753849 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
a=pmft: T38
m=audio 3456 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fntp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cdsc: 4 image udptl t38
```

Step 5:

The terminating Gateway supports T.38, and the RemoteConnectionDescriptor included indicates that the other side supports T.38 as well, so the strict T.38 Call Agent controlled procedure requested can be used. The terminating Gateway supports V.152 and redundancy, and the RemoteConnectionDescriptor included indicates that the other side supports V.152 and redundancy, so Gateway controlled VBD using V.152 and redundancy can be used for modem and text transmissions. The terminating Gateway sends back a success response with its SDP which also includes capability, V.152 and redundancy information:

```
200 2000 OK
I:2

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
a=pmft: T38
m=audio 1296 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fntp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmd:97 vbd=yes
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cdsc: 4 image udptl t38
```

Step 6:

The terminating Call Agent sends back a SIP 200 OK response to the originating Call Agent, which in turn sends a SIP ACK (not shown).

Step 7:

The originating Call Agent in turns sends a ModifyConnection command to the originating Gateway:

```
MDCX 1001 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
M: sendrecv

v=0
o=- 25678 753849 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
a=pmft: T38
m=audio 1296 RTP/AVP 18 96 97
a=rtpmap:96 RED/8000
a=fmtp:96 97/97
a=rtpmap:97 PCMU/8000
a=gpmid:97 vbd=yes
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cdsc: 4 image udptl t38
```

The ModifyConnection command does not repeat the LocalConnectionOptions sent previously. As far as fax handling is concerned, the Gateway therefore attempts to continue using the current fax handling procedure, i.e., strict Call Agent controlled T.38. Since the capability information indicates the other side supports T.38, the Gateway will in fact be able to use the strict Call Agent controlled T.38 procedure. Since the RemoteConnectionDescriptor indicates the other side supports V.152 and redundancy, the Gateway will in fact be able to use the V.152 VBD procedure with redundancy.

Step 8:

The Gateway acknowledges the command. At this point, a call is established using G.729 encoding, and if a fax call is detected, the Call Agent controlled T.38 procedure will be initiated. If a modem or text call is detected, the V.152 VBD procedure will be initiated.

Step 9-10:

The terminating Gateway detects the T.30 CED tone (a.k.a. V.25 ANS). Since both fax and modem calls can start with this sequence, it is not possible to determine that this is a fax call until step 14, where the V.21 fax preamble is detected. The terminating Gateway begins transmitting RTP packets with the negotiated redundant VBD payload type (96).

The "gwvbd(start)" event occurs and is notified to the Call Agent:

```
NTFY 2500 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: vbd/gwvbd(start, rc=ANS, codec=audio/RED, coord=v152ptsw)
X: 20
```

Step 11:

The Call Agent acknowledges the Notify command:

```
200 2500 OK
```

Step 12:

Upon receiving a RTP packet with the redundant VBD payload type (96), the originating Gateway begins transmitting RTP packets with the redundant VBD payload type.

The "gwvbd(start)" event occurs and is notified to the Call Agent:

```
NTFY 1500 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
O: vbd/gwvbd(start, rc=PTSW, codec=audio/RED)
X: 1
```

Step 13:

The Call Agent acknowledges the Notify command:

```
200 1500 OK
```

Step 14 - 15:

The terminating Gateway detects the V.21 fax preamble.

The terminating Gateway is using the Call Agent controlled T.38 strict procedure for fax calls, so the "t38(start)" event occurs:

```
NTFY 2500 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
O: fxr/t38(start)
X: 20
```

Step 16:

The Call Agent acknowledges the Notify command:

```
200 2500 OK
```

Step 17:

The Call Agent then instructs the terminating Gateway to change to using the "image/t38" MIME type instead:

```
MDCX 2002 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2
```

```
L: a:image/t38
R: fxr/t38
X: 21
```

Note that the Call Agent is no longer requesting notification of the "gwvbd" event.

Step 18:

The terminating Gateway sends back a success response with its SDP which also includes the "image/t38" media description:

```
200 2002 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
m=image 1296 udptl t38
a=sqn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cpar: a=rtpmap:96 RED/8000
a=cpar: a=fmtp:96 97/97
a=cpar: a=rtpmap:97 PCMU/8000
a=cpar: a=gpmd:97 vbd=yes
a=cdsc: 4 image udptl t38
```

The gwvbd procedure ends due to the media type change. The "gwvbd(stop)" event would normally be notified at this point, however, the Call Agent is no longer requesting notification of the "gwvbd" event. The Call Agent would have inferred from the "t38(start)" event that the gwvbd procedure ended.

Step 19:

The terminating Call Agent sends a re-INVITE to the originating Call Agent with the updated SDP.

Step 20:

The originating Call Agent then sends a ModifyConnection command to the originating Gateway:

```
MDCX 1003 ds/dsl-1/1@gw-o.whatever.net MGCP 1.0
C: 1
I: 1
R: fxr/t38
X: 2

v=0
o=- 25678 753850 IN IP4 128.96.41.2
s=-
c=IN IP4 128.96.41.2
t=0 0
```



```
m=image 1296 udptl t38
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cpar: a=rtpmap:96 RED/8000
a=cpar: a=fmtp:96 97/97
a=cpar: a=rtpmap:97 PCMU/8000
a=cpar: a=gpmd:97 vbd=yes
a=cdsc: 4 image udptl t38
```

Step 21:

The originating Gateway changes to T.38 and sends back a success response with updated SDP:

```
200 1003 OK

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cpar: a=rtpmap:96 RED/8000
a=cpar: a=fmtp:96 97/97
a=cpar: a=rtpmap:97 PCMU/8000
a=cpar: a=gpmd:97 vbd=yes
a=cdsc: 4 image udptl t38
```

Again, the gwvbd procedure ends due to the media type change. The "gwvbd(stop)" event would normally be notified at this point, however, the Call Agent is no longer requesting notification of the "gwvbd" event.

Step 22:

The originating Call Agent sends a SIP 200 OK response with the updated SDP to the terminating Call Agent, which in turn sends a SIP ACK (not shown).

Step 23:

The terminating Call Agent sends a ModifyConnection with the updated SDP to the terminating Gateway:

```
MDCX 2002 ds/dsl-1/2@gw-t.whatever.net MGCP 1.0
C: 2
I: 2

v=0
o=- 25678 753850 IN IP4 128.96.41.1
s=-
c=IN IP4 128.96.41.1
t=0 0
m=image 3456 udptl t38
```

```

a=sgn: 0
a=cdsc: 1 audio RTP/AVP 18 96 97
a=cpar: a=rtpmap:96 RED/8000
a=cpar: a=fmtp:96 97/97
a=cpar: a=rtpmap:97 PCMU/8000
a=cpar: a=gpmd:97 vbd=yes
a=cdsc: 4 image udptl t38

```

Step 24 - 29:

These steps correspond to the Call Agent controlled T.38 strict procedure as defined in the MGCP Fax (FXR) package [FAXPKG].

8.

Security Considerations

The MGCP VBD package itself is not known to introduce any new security concerns. VBD media is transported over RTP. If security services are in place to protect RTP media streams, these will also be in effect for the VBD media stream.

The MGCP GPMD package does not require any further security considerations beyond those indicated in MGCP [RFC3435].

9.

IANA Considerations

The IANA is hereby requested to register the following MGCP packages,

Package Title	Name	Version
-----	----	-----
Voiceband Data	VBD	0
General-Purpose Media Descriptor Parameter	GPMD	0

and the following LocalConnectionOption,

Field	Name
-----	----
General-Purpose Media Descriptor Parameter	gpmd

10.

Summary of Changes

None (at this time).

11.

Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3550] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 3550, July 2003.

[RFC3435] F. Andreassen, B. Foster, "Media Gateway Control Protocol (MGCP) Version 1.0", RFC 3435, January 2003.

[RFC2198] C. Perkins, I. Kouvelas, O. Hodson, V. Hardman, M. Handley, J.C. Bolot, A. Vega-Garcia, S. Fosse-Parisis, "RTP Payload for Redundant Audio Data", RFC 2198, September 1997.

[RFC2327] M. Handley, V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

[V1501] ITU-T Recommendation V.150.1, "Modem-over-IP networks: Procedures for the end-to-end connection of V-series DCEs", 01/2003.

[V152] ITU-T Recommendation V.152, "Procedures for supporting Voice-Band Data over IP Networks", 01/2005.

[RFC2833] H. Schulzrinne, S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, May 2000.

[FAXPKG] F. Andreassen, "Media Gateway Control Protocol Fax Package", draft-andreasen-mgcp-fax-04.txt, March 23, 2005.

[FMPKG] B. Foster, F. Andreassen, "Basic Media Gateway Control Protocol (MGCP) Packages", RFC 3660, December 2003.

[H2482] ITU-T Recommendation H.248.2, "Gateway control protocol: Facsimile, text conversation and call discrimination packages", 11/2000.

[H1501A1] ITU-T Recommendation V.150.1 Amendment 1, "Modem-over-IP networks: Procedures for the end-to-end connection of V-series DCEs, Amendment 1: Modification to SSE reason identifier codes to support voice band data and text relay", 01/2005.

[RFC2234] D. Crocker, P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.

[RFC2733] J. Rosenberg, H. Schulzrinne, "An RTP Payload Format for Generic Forward Error Correction", RFC 2733, December 1999.

12.

Informative References

[RFC3261] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[T38] ITU-T Recommendation T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", 04/2004.

[RFC2793] G. Hellstrom, "RTP Payload for Text Conversation", RFC 2793, May 2000.

13.

Acknowledgements

Several people have contributed to the development of the MGCP VBD and GPMD packages and the use of the MIME subtypes "RED" and "parityfec" with the FM package for VBD with redundancy and FEC. In particular, the authors would like to thank Flemming Andreasen, John Atkinson, Bill Foster and the CableLabs PacketCable TGCP/NCS focus team for their contributions.

14.

Author's Address

Joe Stone
Cisco Systems
2200 East President George Bush Highway
Richardson, Texas 75082

Email: joestone@cisco.com

Rajesh Kumar
Cisco Systems
771 Alder Drive
Milpitas, California 95035

Email: rkumar@cisco.com

Appendix B. Connection Mode

An MGCP connection can establish one or more media streams. These streams are either incoming or outgoing. The "connection mode" parameter controls the flow of media on the media stream. When there is only one connection to an endpoint, the mapping of these streams is straightforward. 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 mix differently with each other and the streams going to/from the handset. Table 20 below describes how media from different connections should be mixed when one or more connections exist. The table assumes that there are no signals being applied on a connection. Table 20 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 whatsoever

Table 20 - Control of Media Streams by Connection Mode

		Connection A Mode						
		sendonly	recvonly	sendrecv	confrnce	inactive	netwloop/ netwtest	replcate
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						A _{out} = A _{in} B _{out} = B _{in} H _{out} =NA	A _{out} = H _{in} B _{out} = B _{in} H _{out} =NA
	replcate							A _{out} = H _{in} B _{out} = H _{in} H _{out} =NA

For clarity, Table 20 above is repeated below in graphical form:




Figure 6 - Control of Media Streams by Connection Mode - Graphical Representation

If there are three or more connections, their media will still be mixed as defined in the Table 20 above. If internal resources are not available such that the media can not be mixed, the gateway should return error code 502 (insufficient resources).

These connections can be graphically represented as such:




Figure 7 - Graphical Representation of Connections

For example, if Connection A is sendrcv, Connection B is confrnc, and Connection C is rcvonly, from Table 20 above, the outputs in each mode will be illustrated in Figure 8 below:




Figure 8 - Individual Media Streams on Connections

A to B to C interaction:

$$A_{out} = H_{in}$$

$$B_{out} = H_{in}$$

$$C_{out} = NA$$

$$H_{out} = B_{in} + A_{in} + C_{in}$$




Figure 9 - A to B to C Interaction

Appendix C. Dynamic Quality-of-Service

In this appendix, we provide additional detail on the usage of Dynamic Quality-of-Service (D-QoS) in NCS. We describe the expected MTA behavior in more detail and include a state machine and pseudo-code that the MTA MUST follow to support the D-QoS behavior described. The PacketCable Dynamic Quality-of-Service Specification [12] should be consulted for further details.

C.1 Introduction

MTAs implementing support for Dynamic Quality of Service need 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.

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 deleted.

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 identified in Table 21 below for the resource reservation and commit:

Table 21 - Classifiers for Resource Reservation and Commit: Remote Connection Descriptor Provided

		MTA-o (DOCSIS)
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(SDP-t)** refers to the media port in MTA-t'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 classifiers identified in Table 22 MUST be used:

Table 22 - Classifiers for Resource Reservation and Commit: Remote Connection Descriptor Not Provided

		MTA-o (DOCSIS)
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.

C.2 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 (including changing the MIME Media Type)
 - Ptime changes
 - etc.

- The **D-QoS** rules provided in Section 7.3.3.

As a result of a CreateConnection or a ModifyConnection command, if the DOCSIS QoS parameters, Classifier or Authorization Block changed from the previous state, then a new DSA/DSC message **MUST** be sent by the E-MTA to the CMTS containing the updated parameters. An E-MTA **MAY** send a DSC message to the CMTS if the DOCSIS QoS parameters, Classifier and Authorization Block have not changed as a result of a CreateConnection or a ModifyConnection command. 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 23 below:

Table 23 - States Related to Connection Modes and DQoS Reservation Parameters

	SendState	RecvState
No Reserve/Commit parameter supplied – connection mode:		
inactive	R	R
sendonly, replcate	C	R
recvonly	R	C
sendrecv, confrnce, netwloop, netwtest	C	C
Reserve/Commit parameter supplied:		
sendresv	R	N, R*
recvresv	N, R*	R
snrcresv	R	R
sendcomt	C	N, R*
recvcomt	N, R*	C
snrccomt	C	C

* 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 10 and Figure 11 below:

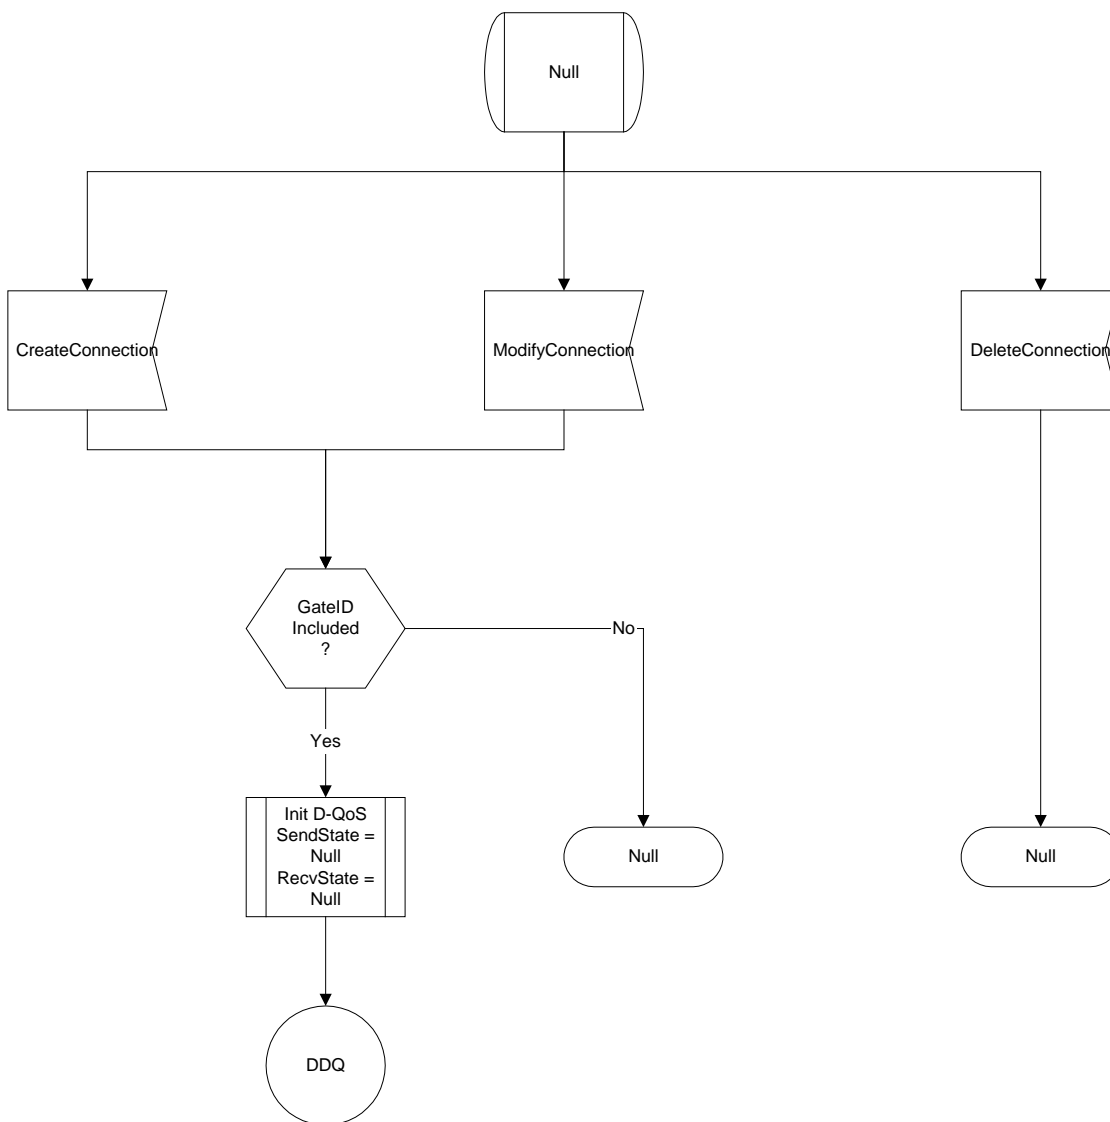
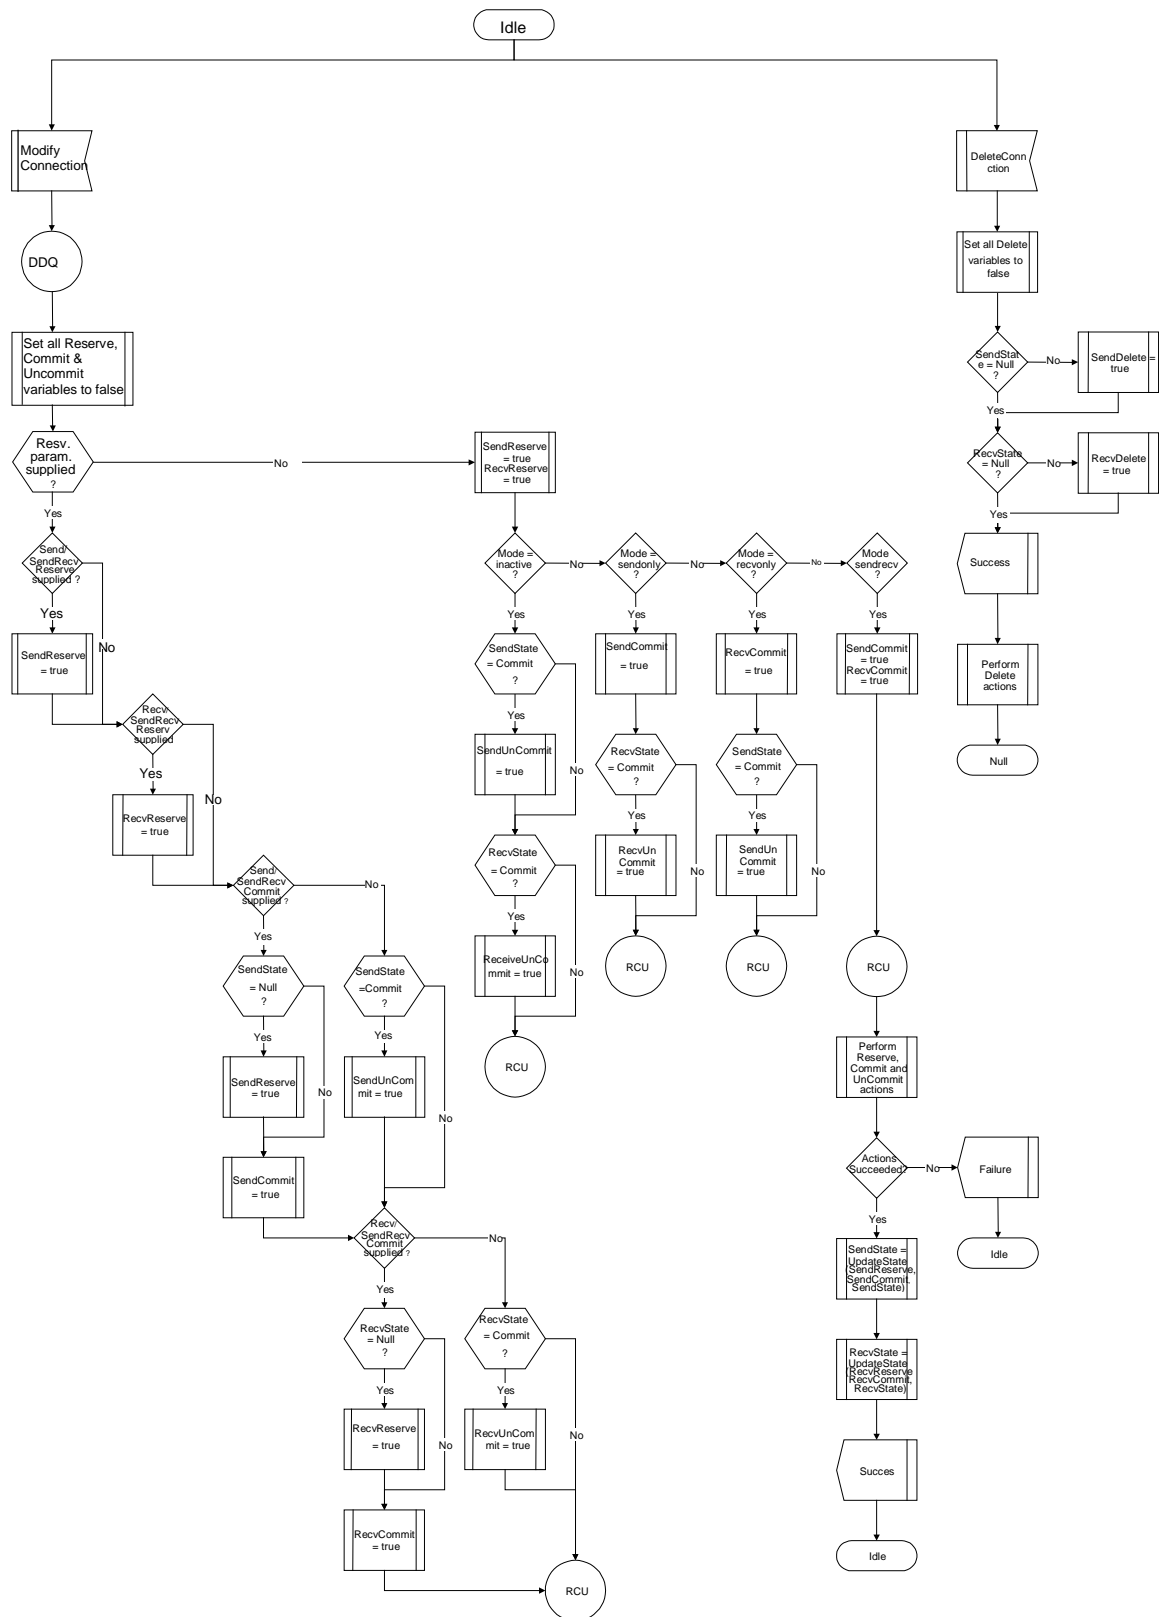


Figure 10 - NCS/D-QoS State Diagram (1:2)



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
    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 RR(*) -- Use wildcard classifier
}

```

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 (DOCSIS QoS)
            SC(RemoteConnectionDescriptor)
    } 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
            RR(*) -- Use wildcard classifier
    }
}

```

```
    }
    If <RemoteConnectionDescriptor provided> then
        RC(RemoteConnectionDescriptor)
    else
        RC(*) -- Use wildcard classifier
}

SendReserveDelete()
If <send resources reserved> then
    SD() -- delete the reservation

ReceiveReserveDelete()
If <receive resources reserved> then
    RD() -- delete the reservation

SendUnCommit()
If <send resources committed> 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
```

Appendix D. Example Command Encodings

This appendix provides examples of commands and responses shown with the actual encoding used. Examples are provided for each command. All commentary shown in the commands and responses is optional.

D.1 NotificationRequest

The first example illustrates a NotificationRequest that will ring a phone and look for an off-hook event:

```
RQNT 1201 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AC
R: hd(N)
S: rg
```

The response indicates that the transaction was successful:

```
200 1201 OK
```

The second example illustrates a NotificationRequest that will look for and accumulate an off-hook event, and then provide dial-tone and accumulate digits according to the digit map provided. The "notified entity" is set to "ca@cal.whatever.net:5678", and since the SignalRequests parameter is empty³⁵, all currently active TO signals will be stopped. All events in the quarantine buffer will be processed, and the list of events to detect in the "notification" and "lockstep" state will include fax tones in addition to the "requested events" and persistent events:

```
RQNT 1202 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AC
R: hd(A, E(S(dl), R(oc, hu, [0-9#*T](D))))
D: (0T|00T|xxxxxxxx|*xx|91xxxxxxxxxxxx|9011x.T)
S:
Q: process
T: ft
```

The response indicates that the transaction was successful:

```
200 1202 OK
```

D.2 Notify

The example below illustrates a Notify message that notifies an off-hook event followed by a 12-digit number beginning with "91". A request identifier correlating the Notify with the NotificationRequest it results from is included. The command is sent to the current "notified entity", which typically will be the actual value supplied in the NotifiedEntity parameter, i.e., "ca@cal.whatever.net:5678" – a failover situation could have changed this:

```
NTFY 2002 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AC
O: hd,9,1,2,0,1,8,2,9,4,2,6,6
```

The Notify response indicates that the transaction was successful:

```
200 2002 OK
```

D.3 CreateConnection

The first example illustrates a CreateConnection command to create a connection on the endpoint specified. The connection will be part of the specified CallId. The LocalConnectionOptions specify that G.711 u-law will be the codec used and the packetization period will be 10 ms. The connection mode will be "receive only":

```
CRCX 1204 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
```

³⁵ It could have been omitted as well.


```
L: p:10, a:PCMU
M: recvonly
```

The response indicates that the transaction was successful, and a connection identifier for the newly created connection is therefore included. A session description for the new connection is included as well – note that it is preceded by an empty line.

```
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 0
a=mptime:10
```

The second example illustrates a CreateConnection command containing a notification request and a RemoteConnectionDescriptor:

```
CRCX 1205 aaln/1@rgw-2569.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
L: p:10, a:PCMU
M: sendrecv
X: 0123456789AD
R: hd
S: rg

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
a=mptime:10
```

The response indicates that the transaction failed, because the phone was already off-hook. Consequently, neither a connection-id nor a session description is returned:

```
401 1205 Phone off-hook
```

Our third example illustrates the use of the provisional response and the three-way handshake. We create another connection this time using dynamic quality of service and acknowledging the previous response received:

```
CRCX 1206 aaln/1@rgw-2569.whatever.net MGCP 1.0 NCS 1.0
K: 1205
C: A3C47F21456789F0
L: p:10, a:PCMU, dq-gi:A735C2
M: inactive

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 18
a=mptime:10 10
```

A provisional response is returned initially:

```
100 1206 Pending
I: DFE233D1

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
```

```

S=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 3456 RTP/AVP 0
a=mptime:10

```

Note that the endpoint elected to support only the PCMU codec, i.e., payload number 0.

A little later, the final response is received:

```

200 1206 OK
K:
DQ-RI: A12D5F1
I: DFE233D1

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
S=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 3456 RTP/AVP 0
a=mptime:10

```

The Call Agent acknowledges the final response as requested:

```
000 1206
```

and the transaction is complete.

D.4 ModifyConnection

The first example shows a ModifyConnection command that simply sets the connection mode of a connection to "send/receive" – the "notified entity" is set as well:

```

MDCX 1209 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
N: ca@cal.whatever.net
M: sendrecv

```

The response indicates that the transaction was successful:

```
200 1209 OK
```

In the second example, we pass a session description and include a notification request with the ModifyConnection command. The endpoint will start playing ring-back tones to the user until it detects audio on the connection specified for the media start event:

```

MDCX 1210 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
M: recvonly
X: 0123456789AE
R: hu, ma@FDE234C8
S: rt

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
S=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 3456 RTP/AVP 0
a=mptime:10

```

The response indicates that the transaction was successful:

```
200 1206 OK
```

D.5 DeleteConnection (From the Call Agent)

In this example, the Call Agent simply instructs the embedded client to delete the connection FDE234C8 on the endpoint specified:

```
DLCX 1210 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
```

The response indicates success, and that the connection was deleted. Connection parameters for the connection are therefore included as well:

```
250 1210 OK
P: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48,
    PC/RPS=782, PC/ROS=45238, PC/RPL=5, PC/RJI=26
```

D.6 DeleteConnection (From the Embedded Client)

In this example, the embedded client sends a DeleteConnection command to the Call Agent to instruct it that a connection on the specified endpoint has been deleted. The ReasonCode specifies the reason for the deletion, and Connection Parameters for the connection are provided as well:

```
DLCX 1210 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
E: 900 - Hardware error
P: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48,
    PC/RPS=782, PC/ROS=45238, PC/RPL=5, PC/RJI=26
```

The Call Agent sends a success response to the gateway:

```
200 1210 OK
```

D.7 DeleteConnection (Multiple Connections From the Call Agent)

In the first example, the Call Agent instructs the embedded client to delete all connections related to call "A3C47F21456789F0" on the specified endpoint:

```
DLCX 1210 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
```

The response indicates success and that the connection(s) were deleted:

```
250 1210 OK
```

In the second example, the Call Agent instructs the embedded client to delete all connections related to all of the endpoints specified:

```
DLCX 1210 aaln/*@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
```

The response indicates success:

```
250 1210 OK
```

D.8 AuditEndpoint

In the first example, the Call Agent wants to learn what endpoints are present on the embedded client specified, hence the use of the "all of" wild-card for the local portion of the endpoint-name:

```
AUEP 1200 *@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
```

The embedded client indicates success and includes a list of endpoint names:

```
200 1200 OK
Z: aaln/1@rgw-2567.whatever.net
Z: aaln/2@rgw-2567.whatever.net
```

In the second example, the capabilities of one of the endpoints is requested:

```
AUEP 1201 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
F:A
```

The response indicates success and the capabilities as well. Two codecs are supported, however with different capabilities. Consequently two separate capability sets are returned. Again, each capability set is to be returned on a single line. The example below shows multiple lines due to formatting restraints:

```
200 1201 OK
A: a:PCMU,p:10-100,e:on,s:off, v:L;S, m:sendonly;
    recvonly;sendrecv;inactive;netwloop;netwtest
A: a:G729,p:30-90, e:on,s:on,v:L;S, m:sendonly;
    recvonly;sendrecv;inactive;confrnce;netwloop
```

In the third example, the Call Agent audits all possible information for the endpoint:

```
AUEP 2002 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
F: R,D,S,X,N,I,T,O,ES,VS,E,MD
```

The response indicates success:

```
200 2002 OK
R: L/hd,L/hu,oc(N),[0-9](N)
D:
S: vmwi(+)
X: 0123456789B1
N: Call-agent@ca.whatever.net
I: 32F345E2
T: L/hd,L/hu,L/ft
O: hd,9,1,2
ES: hd
VS: MGCP 1.0, MGCP 1.0 NCS 1.0
E: 000
MD: 4000
```

The list of requested events contains three events. Where no package name is specified, the default package is assumed. The same goes for actions, so the default action – Notify - must therefore be assumed for the "L/hu" event. The omission of a value for the "digit map" means the endpoint currently does not have a digit map. There are currently no active time-out signals, however the OO signal "vmvi" is currently on and is consequently included – in this case it was parameterized, however the parameter could have been excluded. The current "notified entity" refers to an IP-address and only a single connection exists for the endpoint. The current value of DetectEvents is "ft", and the list of ObservedEvents contains the four events specified. Finally, the event-states audited reveals that the phone was off-hook at the time the transaction was processed.

D.9 AuditConnection

The first example shows an AuditConnection command where we audit the CallId, NotifiedEntity, LocalConnectionOptions, Connection Mode, LocalConnectionDescriptor, and the Connection Parameters:

```
AUCX 2003 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
I: 32F345E2
F: C,N,L,M,LC,P
```

The response indicates success and includes information for the RequestedInfo:

```

200 2003 OK
C: A3C47F21456789F0
N: ca@cal.whatever.net
L: p:10, a:PCMU
M: sendrecv
P: PS=395, OS=22850, PR=615, OR=30937, PL=7, JI=26, LA=47,
    PC/RPS=615, PC/ROS=30937, PC/RPL=5, PC/RJI=26

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 0
a=mptime:10

```

In the second example, we request to audit RemoteConnectionDescriptor and LocalConnectionDescriptor:

```

AUCX 1203 aaln/2@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
I: FDE234C8
F: RC,LC

```

The response indicates success, and includes information for the RequestedInfo. In this case, no RemoteConnectionDescriptor exists, hence only the protocol version field is included for the RemoteConnectionDescriptor:

```

200 1203 OK

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 0
a=mptime:10

v=0

```

D.10 RestartInProgress

The first example illustrates a RestartInProgress message sent by an embedded client to inform the Call Agent that the specified endpoint will be taken out of service in 300 seconds:

```

RSIP 1200 aaln/1@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
RM: graceful
RD: 300

```

The Call Agent's response indicates that the transaction was successful:

```

200 1200 OK

```

In the second example, the RestartInProgress message sent by the embedded client informs the Call Agent, that all of the embedded client's endpoints are being placed in service in 0 seconds, i.e., they are back in service. The delay could have been omitted as well:

```

RSIP 1204 *@rgw-2567.whatever.net MGCP 1.0 NCS 1.0
RM: restart
RD: 0

```

The Call Agent's response indicates success, and furthermore provides the endpoints in question with a new "notified entity":

```

200 1204 OK
N: CA-1@ca.whatever.net

```

Alternatively, the command could have failed with a new "notified entity" as in:

521 1204 OK
N: CA-1@ca.whatever.net

In that case, the command would then have to be retried (as a new transaction) in order to satisfy the "restart procedure" (see Section 7.4.3.5), this time going to Call Agent "CA-1@whatever.net".

Appendix E. Example Call Flow

In this section we provide an example call flow between two embedded clients, EC-1 and EC-2. It should be noted, that this call flow, although a valid one, is merely an example that may or may not be used in practice.

In the call flow described in Table 24 below, CA refers to the Call Agent, CDB refers to a configuration database, and ACC refers to an accounting database.

Table 24 - Example Call Flow

Usrc-1	EC-1	CA	CDB	ACC	EC-2	Usrc-2
	<-	Notification Request				
	Ack	->				
Off-hook	Notify	->				
	<-	Ack				
(Dial-tone)	<-	Create Connection + Notification Request				
	Ack(SDP1)	->				
Digits	Notify	->				
	<-	Ack				
(progress)	<-	Notification Request				
	Ack	->				
		Query(E.164)	->			
		<-	IP			
		Create Connection(SDP1) + Notification Request	----	----	->	
		<-	----	----	P-Ack(SDP2)	
		<-	----	----	Ack(SDP2)	(ringing)
		Ack	----	----	->	
(ringback)	<-	Modify Connection(SDP2) + Notification Request				
	Ack	->				
		<-	----	----	Notify	Off-hook
		Ack	----	----	->	
	<-	ModifyConnection + Notification Request				
	Ack	->				
	(cut in)	Call start	----	->		
		Notification Request	----	----	->	
		<-	----	----	Ack	
		(Call Established)				
		<-	----	----	Notify	on hook

Usr-1	EC-1	CA	CDB	ACC	EC-2	Usr-2
On-hook	Ack (Perf Data)	Ack	---	---	->	
		<- Delete Connection				
		Delete Connection	---	---	->	
		->				
		<-	---	---	Ack(Perf data)	
		Call end	---	->		
		Notification Request	---	---	->	
		<-	---	---	Ack	
		Notify ->				
		<- Ack				
		<- Notification Request				
		Ack ->				

During these exchanges the NCS profile of MGCP is used by the Call Agent to control both embedded clients. The exchanges occur on two sides.

The first command is a NotificationRequest, sent by the Call Agent to the ingress embedded client. The request will consist of the following lines:

```
RQNT 1201 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AB
R: hd
```

The embedded client, at that point, is instructed to look for an off-hook event, and to report it. It will first send a response to the command, repeating in the response the transaction id that the Call Agent attached to the query and providing a return code indicating success:

```
200 1201 OK
```

When the off hook event is noticed, the embedded client sends a Notify message to the Call Agent:

```
NTFY 2001 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AB
O: hd
```

The Call Agent immediately acknowledges the notification:

```
200 2001 OK
```

The Call Agent examines the services associated to an off hook event for this endpoint (it could take special actions in the case of a direct line, no current subscription, etc.). In most cases, it will send a combined CreateConnection and NotificationRequest command to create a connection, provide dial-tone, and collect DTMF digits³⁶:

```
CRCX 1202 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
```

³⁶ The actual digit map depends on dialing plan in the local area as well as services subscribed to. The digit map presented should be considered an example digit map only.


```

L: p:10, a:PCMU
M: recvonly
N: ca@cal.whatever.net:5678
X: 0123456789AC
R: hu, [0-9#*T](D)
D: (0T | 00T | [2-9]xxxxxx | 1[2-9]xxxxxxxxxx | 011xx.T)
S: dl

```

The embedded client acknowledges the transaction, sending back the identification of the newly created connection and the session description used to receive audio data:

```

200 1202 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 0
a=mptime:10

```

The SDP specification, in our example, specifies the address at which the embedded client is ready to receive audio data (128.96.41.1), the transport protocol (RTP), the RTP port (3456) and the audio profile (AVP). The audio profile refers to 9 [3], which defines that the payload type 0 has been assigned for G.711 u-law transmission (also, see [9]).

The embedded client will start accumulating digits according to the digit map. When a digit map match subsequently occurs, the embedded client will notify the observed events to the Call Agent:

```

NTFY 2002 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
N: ca@cal.whatever.net:5678
X: 0123456789AC
O: 1,2,0,1,8,2,9,4,2,6,6

```

The Call Agent immediately acknowledges that notification:

```

200 2002 OK

```

At this stage, the Call Agent will send a NotificationRequest, to stop collecting digits yet continue to watch for an on-hook transition. The Call Agent furthermore decides to acknowledge receipt of the responses for transaction 1202:

```

RQNT 1203 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
K: 1202
X: 0123456789AD
R: hu

```

The embedded client immediately acknowledges that command:

```

200 1203 OK

```

The Call Agent must now create a connection on the egress embedded client, EC-2, and ring the phone attached to the embedded client as well. It does so by sending a combined CreateConnection and NotificationRequest command to the embedded client:

```

CRCX 2001 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0

```

```

C: A3C47F21456789F0
L: p:10, a:PCMU
M: sendrecv
X: 0123456789B0
R: hd
S: rg

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
a=mptime:10

```

The egress embedded client, at that point, is instructed to ring the phone, and to look for an off-hook event, and report it. The off-hook event and ringing signal are synchronized, so when the off-hook event occurs, ringing will stop. The create connection portion of the command has the same parameters as the command sent to the ingress embedded client, with two differences:

- The endpoint identifier points towards the outgoing circuit,
- The message carries the session description returned by the ingress embedded client,
- Because the session description is present, the "mode" parameter is set to "send/receive".

We observe that the call identifier is identical for the two connections. This is normal since the two connections belong to the same call.

We assume, this command does not finish executing immediately³⁷, and a provisional response is therefore returned by the egress embedded client acknowledging the command, sending in the session description its own parameters such as address, ports and RTP profile as well as the connection identifier for the new connection:

```

100 2001 Pending
I: 32F345E2

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
s=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 1297 RTP/AVP 0
a=mptime:10

```

Once the transaction finishes execution, the embedded client sends the final response to the Call Agent, repeating the information it provided in the provisional response:

```

200 2001 OK
K:
I: 32F345E2

```

³⁷ This could, e.g., be due to external resource reservation, although we did not include that in our example.

```

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
s=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 1297 RTP/AVP 0
a=mptime:10

```

When the Call Agent receives the final response, it notices the presence of the empty Response Acknowledgement attribute and therefore issues a Response Acknowledgement for the transaction:

```
000 2001
```

The Call Agent will relay the information to the ingress embedded client, and instruct it to generate local ringback tones, using a combined ModifyConnection and NotificationRequest command:

```

MDCX 1204 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
M: recvonly
X: 0123456789AE
R: hu
S: rt

```

```

v=0
o=- 4723891 7428910 IN IP4 128.96.63.25
s=-
c=IN IP4 128.96.63.25
t=0 0
m=audio 1297 RTP/AVP 0
a=mptime:10

```

The embedded client immediately acknowledges the modification:

```
200 1204 OK
```

At this stage, the Call Agent has established a half duplex transmission path. The phone attached to the ingress embedded client will be able to receive the signals, such as tones or announcements, that may be generated in case of any errors, as well as the initial speech that most likely will be generated when the egress user answers the phone.

When the off hook event is observed, the egress embedded client sends a Notify message to the Call Agent:

```

NTFY 3001 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789B0
O: hd

```

The call agent immediately acknowledges that notification:

```
200 3001 OK
```

The Call agent now sends a combined ModifyConnection and NotificationRequest to the ingress embedded client, to place the connection in send/receive mode and stop the ringback tones:

```

MDCX 1206 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0

```

```
I: FDE234C8
M: sendrecv
X: 0123456789AF
R: hu
```

The embedded client immediately responds to the command:

```
200 1206 OK
```

In parallel, the Call Agent asks the egress embedded client to notify the occurrence of an on-hook event. It does so by sending a NotificationRequest to the embedded client³⁸:

```
RQNT 2002 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789B1
R: hu
```

The embedded client immediately responds to the command:

```
200 2002 OK
```

At this point, the call is fully established.

At some later point in time, the phone attached to the egress embedded client, in our scenario, goes on-hook. This event is notified to the Call Agent, according to the policy received in the last NotificationRequest by sending a Notify command:

```
NTFY 2003 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789B1
O: hu
```

The Call Agent immediately responds to the command:

```
200 2003 OK
```

The Call Agent now determines that the call is ending, and it therefore sends both embedded clients a DeleteConnection command:

```
DLCX 1207 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: FDE234C8
DLCX 2004 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0
C: A3C47F21456789F0
I: 32F345E2
```

The embedded clients will respond with acknowledgements that include the connection parameters for the connection:

```
250 1207 OK
P: PS=1245, OS=62345, PR=780, OR=45123, PL=10, JI=27, LA=48,
   PC/RPS=790, PC/ROS=45700, PC/RPL=15, PC/RJI=26
250 2004 OK
P: PS=790, OS=45700, PR=1230, OR=61875, PL=15, JI=27, LA=48,
   PC/RPS=1245, PC/ROS=62345, PC/RPL=10, PC/RJI=27
```

³⁸ It should be noted, that although on-hook is a persistent event, lockstep mode requires the Call Agent to send a new NotificationRequest to the embedded client.

The Call Agent will also issue a new NotificationRequest to the egress embedded client, to be ready to receive the next off-hook event detected by the embedded client:

```
RQNT 2005 aaln/1@ec-2.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789B2
R: hd
```

The embedded client will acknowledge this message:

```
200 2005 OK
```

Finally, the ingress embedded client hangs up the phone thereby generating a Notify message to the Call Agent:

```
NTFY 1208 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789AF
O: hu
```

The Call Agent immediately responds to the command:

```
200 1208 OK
```

The Call Agent will then issue a new NotificationRequest to the ingress embedded client, to be ready to receive the next off-hook event detected by the embedded client:

```
RQNT 1209 aaln/1@ec-1.whatever.net MGCP 1.0 NCS 1.0
X: 0123456789B3
R: hd
```

The embedded client will acknowledge this message:

```
200 1209 OK
```

Both embedded clients, at this point, are ready for the next call.

Appendix F. Compatibility Information

This appendix provides NCS protocol compatibility information.

F.1 MGCP Compatibility

NCS is a profile of MGCP 1.0 [1], 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** PacketCable 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** PacketCable 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, PacketCable 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.
- **Packetization Period** A new multiple packetization period LocalConnectionOption has been defined, and the MTA is not allowed to choose codecs with a frame size that is inconsistent with the packetization period(s) specified by the Call Agent. Also, codec negotiation differs by depending on and negotiating packetization periods as well as codecs.

Finally, it should be noted, that NCS provides interpretations of and in some cases additional specification or clarification of the base MGCP protocol behavior that may or may not reflect the intended MGCP behavior.

Appendix G. ABNF Grammar for NCS

RFC 3435 includes a formal description of the MGCP protocol syntax following the "Augmented BNF for Syntax Specifications". This formal description is referenced by developers for the creation of interoperable devices. A copy of the MGCP protocol syntax, annotated and edited to indicate its applicability to PacketCable specifications, is provided in this appendix.

Implementations **SHOULD** conform to the portions of this ABNF grammar that relate to their respective specifications, i.e., NCS, TGCP. Also note that there are a few parameter encodings (e.g., embedded request, digit maps, vendor extension names) where the NCS grammar and/or TGCP grammar differ from the MGCP grammar.

Five annotations are used to distinguish between four different cases:

1. The language of the RFC has been changed to accommodate NCS and TGCP requirements.
2. The language of the RFC is applicable to NCS only (and possibly MGCP).
3. The language of the RFC is applicable to TGCP only (and possibly MGCP).
4. The language of the RFC is only applicable to NCS and TGCP.
5. The language of the RFC is only applicable to MGCP.

The language in each case is indicated by a different typeface, as defined in the change description below.

;RFC 3435 grammar changed to accommodate NCS and TGCP

;Bold indicates NCS only (and possibly MGCP)

;Italics indicates TGCP only (and possibly MGCP)

;Bold italics indicates NCS and TGCP only

;Text in grey is for MGCP only

MGCPMessage = MGCPCommand / MGCPResponse

MGCPCommand = MGCPCommandLine 0*(MGCPParameter) [EOL *SDPinformation]

MGCPCommandLine = MGCPVerb 1*(WSP) transaction-id 1*(WSP)
endpointName 1*(WSP) MGCPversion EOL

MGCPVerb = "EPCF" / "CRCX" / "MDCX" / "DLCX" / "RQNT"
/ "NTFY" / "AUEP" / "AUCX" / "RSIP" / extensionVerb

extensionVerb = ALPHA 3(ALPHA / DIGIT) ; experimental starts with X

transaction-id = 1*9(DIGIT)

endpointName = LocalEndpointName "@" DomainName

LocalEndpointName = LocalNamePart 0*("/") LocalNamePart

LocalNamePart = AnyName / AllName / NameString

AnyName = "\$"

AllName = "*"

NameString = 1*(range-of-allowed-characters)

; VCHAR except "\$", "*", "/", "@"

range-of-allowed-characters = %x21-23 / %x25-29 / %x2B-2E
/ %x30-3F / %x41-7E

DomainName = 1*255(ALPHA / DIGIT / "." / "-") ; as defined
 / "#" number ; in RFC 821
 / "[" IPv4address / IPv6address "]" ; see RFC 2373

; Rewritten to ABNF from RFC 821
 number = 1*DIGIT

;From RFC 2373
 IPv6address = hexpart [":" IPv4address]
 IPv4address = 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT "." 1*3DIGIT
 ; this production, while occurring in RFC2373, is not referenced
 ; IPv6prefix = hexpart "/" 1*2DIGIT
 hexpart = hexseq / hexseq ":" [hexseq] / "::" [hexseq]
 hexseq = hex4 *(":" hex4)
 hex4 = 1*4HEXDIG

MGCPversion = "MGCP" 1*(WSP) 1*(DIGIT) "." 1*(DIGIT)
 [1*(WSP) ProfileName]
 ProfileName = "NCS 1.0" ; For NCS
 / "TGCP 1.0" ; For TGCP
 / VCHAR *(WSP / VCHAR)

MGCPPParameter = ParameterValue EOL

; Check infoCode if more parameter values defined
 ; Most optional values can only be omitted when auditing
 ParameterValue = ("K" ":" 0*(WSP) [ResponseAck])
 / ("B" ":" 0*(WSP) [BearerInformation])
 / ("C" ":" 0*(WSP) CallId)
 / ("I" ":" 0*(WSP) [ConnectionId])
 / ("N" ":" 0*(WSP) [NotifiedEntity])
 / ("X" ":" 0*(WSP) [RequestIdentifier])
 / ("L" ":" 0*(WSP) [LocalConnectionOptions])
 / ("M" ":" 0*(WSP) ConnectionMode)
 / ("R" ":" 0*(WSP) [RequestedEvents])
 / ("S" ":" 0*(WSP) [SignalRequests])
 / ("D" ":" 0*(WSP) [DigitMap]) ; For NCS (and MGCP)
 / ("O" ":" 0*(WSP) [ObservedEvents])
 / ("P" ":" 0*(WSP) [ConnectionParameters])
 / ("E" ":" 0*(WSP) ReasonCode)
 / ("Z" ":" 0*(WSP) [SpecificEndpointID])
 / ("Z2" ":" 0*(WSP) SecondEndpointID)
 / ("I2" ":" 0*(WSP) SecondConnectionID)
 / ("F" ":" 0*(WSP) [RequestedInfo])
 / ("Q" ":" 0*(WSP) QuarantineHandling)
 / ("T" ":" 0*(WSP) [DetectEvents])
 / ("RM" ":" 0*(WSP) RestartMethod)
 / ("RD" ":" 0*(WSP) RestartDelay)
 / ("A" ":" 0*(WSP) [Capabilities])
 / ("ES" ":" 0*(WSP) [EventStates])
 / ("PL" ":" 0*(WSP) [PackageList]) ; Auditing only
 / ("MD" ":" 0*(WSP) MaxMGCPDatagram) ; Auditing only
 / (extensionParameter ":" 0*(WSP) [parameterString])
 / *VersionSupported* ; NCS and TGCP - response only
 / *MaxEndpointIds* ; NCS and TGCP
 / *NumEndpoints* ; NCS and TGCP - response only

<extensionParameter> ":" parameterString defined by NCS and TGCP
*VersionSupported = "VS" ":" MGCPversion *("," 0*(WSP) MGCPversion)*
MaxEndpointIds = "ZM" ":" 0(WSP) 1*16(DIGIT)*
NumEndpoints = "ZN" ":" 0(WSP) 1*16(DIGIT) ; Responses only*

; A final response may include an empty ResponseAck
 ResponseAck = confirmedTransactionIdRange
 ("," 0(WSP) confirmedTransactionIdRange)

confirmedTransactionIdRange = transaction-id ["-" transaction-id]

BearerInformation = BearerAttribute 0*("," 0*(WSP) BearerAttribute)
 BearerAttribute = ("e" ":" BearerEncoding)
 / (BearerExtensionName [":" BearerExtensionValue])
 BearerExtensionName = PackageLCOExtensionName
 BearerExtensionValue = LocalOptionExtensionValue
 BearerEncoding = "A" / "mu"
 CallId = 1*32(HEXDIG)

; The audit request response may include a list of identifiers
 ConnectionId = 1*32(HEXDIG) 0*("," 0*(WSP) 1*32(HEXDIG))
 SecondConnectionId = ConnectionId

NotifiedEntity = [LocalName "@"] DomainName [":" portNumber]
 LocalName = 1*(LocalCallAgentNameCharacters)
 LocalCallAgentNameCharacters = %x21-3F / %x41-7E ;VCHAR except "@"

portNumber = 1*5(DIGIT)

RequestIdIdentifier = 1*32(HEXDIG)

LocalConnectionOptions = LocalOptionValue 0*(WSP)
 0*("," 0*(WSP) LocalOptionValue 0*(WSP))
 LocalOptionValue = ("p" ":" packetizationPeriod)
 / ("a" ":" compressionAlgorithm)
 / ("b" ":" **bandwidth**) ; **Only for capabilities in**
 ; **NCS and TGCP**
 / ("e" ":" echoCancellation)
 / ("gc" ":" gainControl)
 / ("s" ":" silenceSuppression)
 / ("t" ":" typeOfService)
 / ("r" ":" resourceReservation)
 / ("k" ":" encryptiondata)
 / ("nt" ":" (typeOfNetwork /
 supportedTypeOfNetwork))
 / (LocalOptionExtensionName
 [":" LocalOptionExtensionValue])
 / **MPacketizationPeriod** ; **NCS and TGCP only**
 / **RTPciphersuite** ; **NCS and TGCP only**
 / **RTCPciphersuite** ; **NCS and TGCP only**
 / **DQoSGateID** ; **NCS only**
 / **DQoSReservation** ; **NCS only**
 / **DQoSResourceID** ; **NCS only**
 / **DQoSReserveDestination** ; **NCS only**
 / **CallContentId** ; **TGCP only**

/ CallContentDestination ; TGCP only

Capabilities = CapabilityValue 0*(WSP)
 0*("," 0*(WSP) CapabilityValue 0*(WSP))
 CapabilityValue = LocalOptionValue
 / ("v" ":" supportedPackages)
 / ("m" ":" supportedModes)

PackageList = pkgNameAndVers 0*("," pkgNameAndVers)
 pkgNameAndVers = packageName ":" packageVersion
 packageVersion = 1*(DIGIT)

***; For NCS and TGCP, range format is only allowed for capabilities
 ; and not for LocalConnectionOptions.***

packetizationPeriod = 1*4(DIGIT) ["-" 1*4(DIGIT)]
 compressionAlgorithm = algorithmName 0*("," algorithmName)
 algorithmName = 1*(SuitableLCOCharacter)
 bandwidth = 1*4(DIGIT) ["-" 1*4(DIGIT)]
 echoCancellation = "on" / "off"
 gainControl = "auto" / ["-"] 1*4(DIGIT)
 silenceSuppression = "on" / "off"
 typeOfService = 1*2(HEXDIG) ; 1 hex only for capabilities
 resourceReservation = "g" / "cl" / "be"

; encryption parameters are coded as in SDP (RFC 2327)
 encryptiondata = ("clear" ":" encryptionKey)
 / ("base64" ":" encodedEncryptionKey)
 / ("uri" ":" URIToObtainKey)
 / ("prompt") ; defined in SDP, not usable in MGCP!
 encryptionKey = 1*(SuitableLCOCharacter) / quotedString
 ; See RFC 2045
 encodedEncryptionKey = 1*(ALPHA / DIGIT / "+" / "/" / "=")
 URIToObtainKey = 1*(SuitableLCOCharacter) / quotedString

typeOfNetwork = "IN" / "ATM" / "LOCAL" / OtherTypeOfNetwork
 ; Registered with IANA - see RFC 2327
 OtherTypeOfNetwork = 1*(SuitableLCOCharacter)
 supportedTypeOfNetwork = typeOfNetwork *("," typeOfNetwork)
 supportedModes = ConnectionMode 0*("," ConnectionMode)

supportedPackages = packageName 0*("," packageName)

packageName = 1*(ALPHA / DIGIT / HYPHEN) ; Hyphen neither first or last

LocalOptionExtensionName = VendorLCOExtensionName
 / PackageLCOExtensionName
 / OtherLCOExtensionName
 VendorLCOExtensionName = "x" ("+" / "-") 1*32(SuitableExtLCOCharacter)
 PackageLCOExtensionName = packageName "/"
 1*32(SuitablePkgExtLCOCharacter)
 ; must not start with "x-" or "x+"
 OtherLCOExtensionName = 1*32(SuitableExtLCOCharacter)

***; <LocalOptionExtensionName> ":" <LocalOptionExtensionvalue>
 ; defined by NCS/TGCP***

MPacketizationPeriod = "mp" ":" *multiplepacketizationPeriod*

multiplepacketizationPeriod = *mpPeriod* 0*(";" *mpPeriod*)

mpPeriod = 1*4(DIGIT) / HYPHEN

RTPciphersuite = "sc-rtp" ":" *ciphersuite*

RTCPciphersuite = "sc-rtcp" ":" *ciphersuite*

ciphersuite = [AuthenticationAlgorithm] "/" [EncryptionAlgorithm]

AuthenticationAlgorithm = 1*(ALPHA / DIGIT / "-" / "_")

EncryptionAlgorithm = 1*(ALPHA / DIGIT / "-" / "_")

; <LocalOptionExtensionName> ":" <LocalOptionExtensionvalue>

; defined by NCS only

DQoSGateID = "dq-gi" [";" 1*8(HEXDIG)] ; Only empty for
; capabilities

DQoSReservation = "dq-rr" ":" DQoSResMode *(";" DQoSResMode)

DQoSResMode = "sendresv" / "recvresv" / "snrcresv" /
"sendcomt" / "recvcomt" / "snrccomt"

DQoSResourceID = "dq-ri" ":" 1*8(HEXDIG)

DQoSReserveDestination = "dq-rd" ":" IPv4address [";" portNumber]

; <LocalOptionExtensionName> ":" <LocalOptionExtensionvalue>

; defined by TGCP only

CallContentId = "es-cci" ":" 1*8(HEXDIG)

CallContentDestination = "es-ccd" ":" IPv4address ":" portNumber

LocalOptionExtensionValue = (1*(SuitableExtLCOValChar)
/ quotedString)

(";" (1(SuitableExtLCOValChar)
/ quotedString))

;Note: No "data" mode.

ConnectionMode = "sendonly" / "recvonly" / "sendrecv"

/ "confrnce" / "inactive"

/ "loopback" / "conttest" ; TGCP (and MGCP) only

/ "replicate" ; NCS only

/ "netwloop" / "netwtest"

/ ExtensionConnectionMode

ExtensionConnectionMode = PkgExtConnectionMode

PkgExtConnectionMode = packageName "/" 1*(ALPHA / DIGIT)

RequestedEvents = requestedEvent 0*(";" 0*(WSP) requestedEvent)

requestedEvent = (eventName ["(" requestedActions ")"])

/ (eventName "(" requestedActions ")"

"(" eventParameters ")")

eventName = [(packageName / "*") "/"]

(eventId / "all" / eventRange

/ "*" / "#") ; for DTMF

["@" (ConnectionId / "\$" / "*")]

eventId = 1*(ALPHA / DIGIT / HYPHEN) ; Hyphen neither first nor last

eventRange = [" 1*(DigitMapLetter / (DIGIT "-" DIGIT) /

(DTMFLetter "-" DTMFLetter)) "]"

DTMFLetter = "A" / "B" / "C" / "D"

requestedActions = requestedAction 0*(";" 0*(WSP) requestedAction)

requestedAction = "N" / "A"

```

/ "D" ; For NCS (and MGCP)
/ "S" / "I" / "K"
/ "E" "(" EmbeddedRequest ")"
/ ExtensionAction
/ "C" "(" EmbeddedModeChange ; For NCS and TGCP
0*("," 0*WSP EmbeddedModeChange ")" ) ; only
; NCS and TGCP define the Embedded ModifyConnection action.
; MGCP grammar does not allow for the format used in NCS and TGCP:
EmbeddedModeChange = "M" "(" ConnectionMode "(" EmConnectionId ")" ")"
EmConnectionId     = ConnectionId / "$"

```

```

ExtensionAction = PackageExtAction
PackageExtAction = packageName "/" Action [ "(" ActionParameters ")" ]
Action           = 1*ALPHA

```

ActionParameters = eventParameters ; May contain actions

;NOTE: Should tolerate different order when receiving, e.g., for NCS

```

EmbeddedRequest = ( "R" "(" EmbeddedRequestList ")"
[ "," 0*(WSP) "S" "(" EmbeddedSignalRequest ")" ]
[ "," 0*(WSP) "D" "(" EmbeddedDigitMap ")" ] )
/ ( "S" "(" EmbeddedSignalRequest ")"
[ "," 0*(WSP) "D" "(" EmbeddedDigitMap ")" ] )
/ ( "D" "(" EmbeddedDigitMap ")" )
/ NCSTGCPEmbeddedRequest

```

;Text below is for NCS and TGCP only. The difference compared to MGCP is simply that the order of the items is not fixed. Also for TGCP Digit Maps are not used

```

NCSTGCPEmbeddedRequest = NCSTGCPEmbeddedRequestItem
                        *2( "," 0*(WSP) NCSTGCPEmbeddedRequestItem )
NCSTGCPEmbeddedRequestItem = ("R" "(" EmbeddedRequestList ")" )
                        / ("S" "(" EmbeddedSignalRequest ")" )
                        / ("D" "(" EmbeddedDigitMap ")" )

```

```

EmbeddedRequestList = RequestedEvents
EmbeddedSignalRequest = SignalRequests
EmbeddedDigitMap = DigitMap

```

```

SignalRequests = SignalRequest 0*("," 0*(WSP) SignalRequest )
SignalRequest = eventName [ "(" eventParameters ")" ]

```

```

eventParameters = eventParameter 0*("," 0*(WSP) eventParameter)
eventParameter = eventParameterValue
                / eventParameterName "=" eventParameter
                / eventParameterName "(" eventParameters ")"
eventParameterString = 1*(SuitableEventParamCharacter)
eventParameterName = eventParameterString
eventParameterValue = eventParameterString / quotedString

```

```

; For NCS (and MGCP)
DigitMap = DigitString / "(" DigitStringList ")"
DigitStringList = DigitString 0*("|" DigitString )
DigitString = 1*(DigitStringElement)
DigitStringElement = DigitPosition [ "." ]
DigitPosition = DigitMapLetter / DigitMapRange

```

; NOTE "X" is now included

DigitMapLetter = DIGIT / "#" / "*" / "A" / "B" / "C" / "D" / "T"
/ "X" / ExtensionDigitMapLetter

ExtensionDigitMapLetter = "E" / "F" / "G" / "H" / "I" / "J" / "K"
/ "L" / "M" / "N" / "O" / "P" / "Q" / "R"
/ "S" / "U" / "V" / "W" / "Y" / "Z"

; NOTE "[x]" is now allowed in MGCP.

; In NCS, only the "x" form is allowed

DigitMapRange = "[" 1*DigitLetter "]"

/ "X" ; Added for NCS only

DigitLetter = *((DIGIT "-" DIGIT) / DigitMapLetter)

ObservedEvents = SignalRequests

EventStates = SignalRequests

ConnectionParameters = ConnectionParameter

0*("," 0*(WSP) ConnectionParameter)

ConnectionParameter = ("PS" "=" packetsSent)

/ ("OS" "=" octetsSent)

/ ("PR" "=" packetsReceived)

/ ("OR" "=" octetsReceived)

/ ("PL" "=" packetsLost)

/ ("JI" "=" jitter)

/ ("LA" "=" averageLatency)

/ (ConnectionParameterExtensionName

"=" ConnectionParameterExtensionValue)

/ *RemotePacketsSent*

/ *RemoteOctetsSent*

/ *RemotePacketsLost*

/ *RemoteJitter*

; NCS and TGCP define the following four connection parameter extension

; names:

RemotePacketsSent = "PC/RPS" "=" packetsSent

RemoteOctetsSent = "PC/ROS" "=" octetsSent

RemotePacketsLost = "PC/RPL" "=" packetsLost

RemoteJitter = "PC/JI" "=" jitter

packetsSent = 1*9(DIGIT)

octetsSent = 1*9(DIGIT)

packetsReceived = 1*9(DIGIT)

octetsReceived = 1*9(DIGIT)

packetsLost = 1*9(DIGIT)

jitter = 1*9(DIGIT)

averageLatency = 1*9(DIGIT)

ConnectionParameterExtensionName = VendorCPEExtensionName

/ PackageCPEExtensionName

VendorCPEExtensionName = "X" "-" 2*ALPHA

/ *NCSTGCPVendorCPEExtensionName*

;Text below is for NCS and TGCP only. The difference compared to MGCP

;is simply that MGCP requires 2 alpha characters whereas NCS and TGCP

;allow 2 or 3 alpha characters for VendorCPEExtensionName

NCSTGCPVendorCPEExtensionName = "X" "-" 2*3ALPHA

PackageCPEExtensionName = packageName "/" CPName
 CPName = 1*(ALPHA / DIGIT / HYPHEN)
 ConnectionParameterExtensionValue = 1*9(DIGIT)

MaxMGCPDatagram = 1*9(DIGIT)

ReasonCode = 3DIGIT
 [1*(WSP) "/" packageName] ; Only for 8xx
 [WSP 1*(%x20-7E)]

SpecificEndpointID = endpointName
 SecondEndpointID = endpointName

RequestedInfo = infoCode 0*(" 0*(WSP) infoCode)

infoCode = "B" / "C" / "I" / "N" / "X" / "L" / "M" / "R" / "S"
 / "D" ; **For NCS (and MGCP) only**
 / "O" / "P" / "E" / "Z" / "Q" / "T" / "RC" / "LC"
 / "A" / "ES" / "RM" / "RD" / "PL" / "MD" / extensionParameter
 / "VS" / "ZM" / "ZN" ; **NCS and TGCP define these**
 ; **three extensionParameters**

;NCS and TGCP allows for process and loop control in either order

QuarantineHandling = loopControl / processControl
 / (loopControl " 0*(WSP) processControl)
 / (**processControl " 0*(WSP) loopControl**)

loopControl = "step" / "loop"
 processControl = "process" / "discard"

DetectEvents = SignalRequests

RestartMethod = "graceful" / "forced" / "restart" / "disconnected"
 / "cancel-graceful" / extensionRestartMethod
 extensionRestartMethod = PackageExtensionRM
 PackageExtensionRM = packageName "/" 1*32(ALPHA / DIGIT / HYPHEN)
 RestartDelay = 1*6(DIGIT)

extensionParameter = VendorExtensionParameter
 / PackageExtensionParameter
 / OtherExtensionParameter
 VendorExtensionParameter = "X" ("-" / "+") 1*6(ALPHA / DIGIT)
 PackageExtensionParameter = packageName "/"
 1*32(ALPHA / DIGIT / HYPHEN)
 ; must not start with "x-" or "x+"
 OtherExtensionParameter = 1*32(ALPHA / DIGIT / HYPHEN)

;If first character is a double-quote, then it is a quoted-string
 parameterString = (%x21 / %x23-7F) *(%x20-7F) ; first and last must not
 ; be white space
 / quotedString

MGCPResponse = MGCPResponseLine 0*(MGCPPParameter)
 *2(EOL *SDPinformation)

MGCPResponseLine = responseCode 1*(WSP) transaction-id
 [1*(WSP) "/" packageName] ; Only for 8xx
 [WSP responseString] EOL
 responseCode = 3DIGIT
 responseString = *(%x20-7E)

SuitablePkgExtLCOCharacter = SuitableLCOCharacter

SuitableExtLCOCharacter = DIGIT / ALPHA / "+" / "-" / "_" / "&"
 / "!" / "" / "|" / "=" / "#" / "?"
 / "." / "\$" / "*" / "@" / "[" / "]"
 / "^" / "`" / "{" / "}" / "~"

SuitableLCOCharacter = SuitableExtLCOCharacter / "/"

SuitableExtLCOValChar = SuitableLCOCharacter / ":"

; VCHAR except "(", ")", ",", and "="

SuitableEventParamCharacter = %x21 / %x23-27 / %x2A-2B
 / %x2D-3C / %x3E-7E

; NOTE: UTF8 encoded

quotedString = DQUOTE 0*(quoteEscape / quoteChar) DQUOTE
 quoteEscape = DQUOTE DQUOTE
 quoteChar = (%x00-21 / %x23-FF)

EOL = CRLF / LF

HYPHEN = "-"

; See RFC 2327 for proper SDP grammar instead.

SDPinformation = SDPLine CRLF *(SDPLine CRLF) ; see RFC 2327
 SDPLine = 1*(%x01-09 / %x0B / %x0C / %x0E-FF) ; for proper def.

Appendix H. Acknowledgements

This specification was developed and influenced by numerous individuals representing many different vendors and organizations. PacketCable hereby wishes to thank everybody who participated directly or indirectly in this effort. In particular, PacketCable wants to recognize the authors of the original SGCP, MGCP, and IPDC specifications for the work they have done. Furthermore, PacketCable wishes to thank the following individuals for their involvement and contributions to this specification:

Flemming Andreasen (Cisco -Primary Author)
Brian Bailey (3Com)
Burcak Beser (Pacific Broadband)
David Bukovinsky (Blue Sky)
Rex Coldren (AG Communications)
Graeme Currie (Nortel)
Chad Griffiths (Broadcom)
Jack Fijolek (3Com)
Bill Foster (Cisco)
Billy Hare (Arris)
Barry Hoffner (Telcordia)
Christian Huitema (Telcordia)
Keith Kelly (Netspeak)
Satish Kumar (Texas Instruments) David Lide, Dan Thomas (Telogy) Bill Marshall (AT&T)
Michael Mannette (3Com)
Jiri Matousek (Arris)
Sasha Medvinsky (Motorola)
Steve Michelson (AT&T)
Debbie Ng (Nortel)
David Oran (Cisco)
Jeff Orwick (NetSpeak)
Matt Osman (CableLabs)
Daniel Paul (Arris)
John Pickens (Com21)
David Reiter (Motorola)
Don Stanwyck (Lucent)
Bob Stein (Motorola)
Kurt Steinbrenner (3Com)
Dan Thomas (Telogy)
Michael Thomas (Cisco)
Raymond Yeung (Pacific Broadband)
Venkatesh Sunkad (CableLabs)
David Hancock (CableLabs)
Sandeep Sharma (CableLabs)

Joey Padden (CableLabs)

Appendix I. Revision History

The following ECN has been incorporated into PKT-SP-NCS1.5-I02-050812.

ECN	ECN Date	Summary
NCS1.5-N-05.0225-5	3/14/05	Media security parameter usage for T38

The following ECNS have been incorporated into PKT-SP-NCS1.5-I03-070412.

ECN	ECN Date	Summary
NCS1.5-N-0301-2	3/27/06	V.152 Support
NCS1.5-N-07.0385-3	2/26/07	Make the support for packages related to V.152 optional in NCS
NCS1.5-N-07.0402-3	3/12/07	Fix errors in ECN NCS1.5-N-05.0301-2.
NCS1.5-N-07.0409-2	3/12/07	Allow inclusion of t38MaxBitRate attribute in SDP
NCS1.5-N-07.0403-5	3/19/07	Update and correct references in NCS

The following ECN has been incorporated into PKT-SP-NCS1.5-I04-120412.

ECN	ECN Date	Summary
NCS1.5-N-11.0667-2	2/13/12	NCS references to CODEC corrections
