



***Society of Cable
Telecommunications
Engineers***

**ENGINEERING COMMITTEE
Data Standards Subcommittee**

AMERICAN NATIONAL STANDARD

ANSI/SCTE 22-1 2002R2007

**Data-Over-Cable Service Interface Specification DOCSIS
1.0 Radio Frequency Interface (RFI)**

NOTICE

The Society of Cable Telecommunications Engineers (SCTE) Standards are intended to serve the public interest by providing specifications, test methods and procedures that promote uniformity of product, interchangeability and ultimately the long term reliability of broadband communications facilities. These documents shall not in any way preclude any member or nonmember of SCTE from manufacturing or selling products not conforming to such documents, nor shall the existence of such standards preclude their voluntary use by those other than SCTE members, whether used domestically or internationally.

SCTE assumes no obligations or liability whatsoever to any party who may adopt the Standards. Such adopting party assumes all risks associated with adoption of these Standards or Recommended Practices, and accepts full responsibility for any damage and/or claims arising from the adoption of such Standards or Recommended Practices.

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. SCTE shall not be responsible for identifying patents for which a license may be required or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

Patent holders who believe that they hold patents which are essential to the implementation of this standard have been requested to provide information about those patents and any related licensing terms and conditions. Any such declarations made before or after publication of this document are available on the SCTE web site at <http://www.scte.org>.

All Rights Reserved

© Society of Cable Telecommunications Engineers, Inc. 2002, 2007
140 Philips Road
Exton, PA 19341

Table of Contents

1	SCOPE AND PURPOSE	1
1.1	Scope	1
1.2	Requirements	1
1.3	Background	2
1.3.1	Service Goals	2
1.3.2	Reference Architecture	2
1.3.3	Server Location	6
2	FUNCTIONAL ASSUMPTIONS	7
2.1	Broadband Access Network	7
2.2	Equipment Assumptions	7
2.2.1	Frequency Plan	7
2.2.2	Compatibility with Other Services	7
2.2.3	Fault Isolation Impact on Other Users	7
2.3	RF Channel Assumptions	8
2.3.1	Transmission Downstream	8
2.3.2	Transmission Upstream	8
2.4	Transmission Levels	9
2.5	Frequency Inversion	9
3	COMMUNICATION PROTOCOLS	10
3.1	Protocol Stack	10
3.1.1	CM and CMTS as Hosts	10
3.1.2	Data Forwarding Through the CM and CMTS	11
3.2	The MAC Forwarder	13
3.2.1	Example Rules for Data-Link-Layer Forwarding	14
3.3	Network Layer	15
3.4	Above the Network Layer	15
3.5	Data Link Layer	15
3.5.1	LLC Sublayer	15
3.5.2	Link-Layer Security Sublayer	15
3.5.3	MAC Sublayer	15
3.6	Physical Layer	16
3.6.1	Downstream Transmission Convergence Sublayer	16
3.6.2	PMD Sublayer	16
4	PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION	18
4.1	Scope	18
4.2	Upstream	18
4.2.1	Overview	18
4.2.2	Modulation Formats	19
4.2.3	FEC Encode	21
4.2.4	Scrambler (Randomizer)	22
4.2.5	Preamble Prepend	22
4.2.6	Transmit Pre-Equalizer	23
4.2.7	Burst Profiles	24
4.2.8	Burst Timing Convention	26
4.2.9	Transmit Power Requirements	27
4.2.10	Fidelity Requirements	28
4.2.11	Frame Structure	31
4.2.12	Signal Processing Requirements	32
4.2.13	Upstream Demodulator Input Power Characteristics	33
4.2.14	Upstream Electrical Output from the CM	33

4.3	Downstream	34
4.3.1	Downstream Protocol	34
4.3.2	Scalable Interleaving to Support Low Latency.....	34
4.3.3	Downstream Frequency Plan.....	34
4.3.4	CMTS Output Electrical.....	34
4.3.5	Downstream Electrical Input to CM.....	35
4.3.6	CM BER Performance.....	36
4.3.7	CMTS Timestamp Jitter.....	37
5.	DOWNSTREAM TRANSMISSION CONVERGENCE SUBLAYER.....	38
5.1	Introduction	38
5.2	MPEG Packet Format.....	38
5.3	MPEG Header for DOCSIS Data-Over-Cable	39
5.4	MPEG Payload for DOCSIS Data-Over-Cable.....	39
5.5	Interaction with the MAC Sublayer	40
5.6	Interaction with the Physical Layer.....	41
5.7	MPEG Header Synchronization and Recovery	41
6	MEDIA ACCESS CONTROL SPECIFICATION	42
6.1	Introduction	42
6.1.1	Overview	42
6.1.2	Definitions	42
6.1.3	Future Use.....	43
6.2	MAC Frame Formats.....	43
6.2.1	Generic MAC Frame Format.....	43
6.2.2	Packet-Based MAC Frames.....	47
6.2.3	ATM Cell MAC Frames.....	48
6.2.4	Reserved PDU MAC Frames.....	49
6.2.5	MAC-Specific Headers.....	50
6.2.6	Extended MAC Headers.....	53
6.2.7	Error-Handling.....	55
6.3	MAC Management Messages	56
6.3.1	MAC Management Message Header.....	56
6.3.2	MAC Management Messages.....	57
6.4	Upstream Bandwidth Allocation.....	75
6.4.1	The Allocation Map MAC Management Message.....	76
6.4.2	Map Transmission and Timing.....	79
6.4.3	Protocol Example.....	79
6.4.4	Contention Resolution.....	80
6.4.5	CM Behavior.....	82
6.4.6	Support for Multiple Channels.....	82
6.4.7	Classes of Service.....	83
6.5	Timing and Synchronization	84
6.5.1	Global Timing Reference.....	84
6.5.2	CM Channel Acquisition.....	84
6.5.3	Ranging.....	84
6.5.4	Timing Units and Relationships.....	85
6.6	Data Link Encryption Support	86
6.6.1	MAC Messages.....	86
6.6.2	Framing.....	86
7.	CABLE MODEM - CMTS INTERACTION.....	87
7.1	CMTS Initialization.....	87
7.2	Cable Modem Initialization	87
7.2.1	Scanning and Synchronization to Downstream.....	89
7.2.2	Obtain Upstream Parameters.....	89

7.2.3	Message Flows During Scanning and Upstream Parameter Acquisition	92
7.2.4	Ranging and Automatic Adjustments	93
7.2.5	Establish IP Connectivity	97
7.2.6	Establish Time of Day	98
7.2.7	Transfer Operational Parameters	98
7.2.8	Registration.....	98
7.2.9	Baseline Privacy Initialization	102
7.2.10	Service IDs During CM Initialization.....	102
7.2.11	Multiple-Channel Support	102
7.2.12	Remote RF Signal Level Adjustment	102
7.2.13	Changing Upstream Burst Parameters	104
7.2.14	Changing Upstream Channels	104
7.2.15	Fault Detection and Recovery	106
7.2.16	Prevention of Unauthorized Transmissions.....	107
8	SUPPORTING FUTURE NEW CABLE MODEM CAPABILITIES.....	108
8.1	Setting Up Communications on an Enhanced Basis.....	108
8.1.1	Upstream Enhanced / Downstream Standard	108
8.1.2	Downstream Enhanced / Upstream Enhanced or Standard	108
8.2	Downloading Cable Modem Operating Software.....	108
9	PROVISION FOR OTHER FUTURE CAPABILITIES	110
9.1	Anticipated Physical-Layer Changes.....	110
9.1.1	Adding Upstream Channel and Burst Configuration Settings	110
9.1.2	Downstream Channel Improvements.....	111
9.2	New Network Service Requirements.....	111
9.2.1	Multicast Service IDs	111
9.2.2	RSVP Support for Upstream Traffic	112
9.3	PID Filtering Capability	113
	APPENDIX A. WELL-KNOWN ADDRESSES	114
	APPENDIX B. PARAMETERS AND CONSTANTS	115
	APPENDIX C. CM CONFIGURATION INTERFACE SPECIFICATION.....	116
	APPENDIX D. MAC SUBLAYER SERVICE DEFINITION	133
	APPENDIX E. EXAMPLE BURST PROFILES	136
	APPENDIX F. UPSTREAM MODULATION RATES.....	149
	APPENDIX G. EXAMPLE: MULTIPLE UPSTREAM CHANNELS	154
	APPENDIX H. THE DATA-OVER-CABLE SPANNING TREE PROTOCOL.....	155
	APPENDIX I. ERROR CODES AND MESSAGES	158
	APPENDIX J. REFERENCES	161
	APPENDIX K. GLOSSARY	164
	APPENDIX L. REVISION HISTORY	172
	APPENDIX M. ACKNOWLEDGMENTS.....	173

List of Figures

Figure 1-1.	Transparent IP Traffic Through the Data-Over-Cable System.....	2
Figure 1-2.	Data-Over-Cable Reference Architecture.....	3
Figure 1-3.	Server Location Not At CMTS.....	6
Figure 3-1.	Protocol Stack on the RF Interface.....	10
Figure 3-2.	Data Forwarding Through the CM and CMTS.....	11
Figure 3-3.	Example Condition for Network Loops.....	12
Figure 3-4.	MAC Forwarder.....	14
Figure 4-1.	QPSK Symbol Mapping.....	19
Figure 4-2.	16QAM Gray-Coded Symbol Mapping.....	20
Figure 4-3.	16QAM Differential-Coded Symbol Mapping.....	20
Figure 4-4.	Scrambler Structure.....	22
Figure 4-5.	Transmit Pre-Equalizer Structure.....	23
Figure 4-6.	Nominal Burst Timing.....	26
Figure 4-7.	Worst-Case Burst Timing.....	27
Figure 4-8.	Example Frame Structures with Flexible Burst Length Mode.....	31
Figure 4-9.	Signal-Processing Sequence.....	32
Figure 4-10.	TDMA Upstream Transmission Processing.....	33
Figure 5-1.	Example of Interleaving MPEG Packets in Downstream.....	38
Figure 5-2.	Format of an MPEG Packet.....	38
Figure 5-3.	Packet Format Where a MAC Frame Immediately Follows the pointer_field.....	40
Figure 5-4.	Packet Format with MAC Frame Preceded by Stuffing Bytes.....	40
Figure 5-5.	Packet Format Showing Multiple MAC Frames in a Single Packet.....	40
Figure 5-6.	Packet Format Where a MAC Frame Spans Multiple Packets.....	40
Figure 6-1.	Generic MAC Frame Format.....	44
Figure 6-2.	Upstream MAC/PMD Convergence.....	44
Figure 6-3.	MAC Header Format.....	45
Figure 6-4.	Ethernet/802.3 Packet PDU Format.....	47
Figure 6-5.	ATM Cell MAC Frame Format.....	48
Figure 6-6.	Reserved PDU Format.....	49
Figure 6-7.	Timing MAC Header.....	50
Figure 6-8.	Management MAC Header.....	51
Figure 6-9.	Request MAC Header Format.....	52
Figure 6-10.	Concatenation of Multiple MAC Frames.....	53
Figure 6-11.	Concatenation MAC Header Format.....	53
Figure 6-12.	Extended MAC Format.....	54
Figure 6-13.	MAC Header and MAC Management Message Header Fields.....	56
Figure 6-14.	Format of Packet PDU Following the Timing Header.....	58
Figure 6-15.	Upstream Channel Descriptor.....	58
Figure 6-16.	Top-Level Encoding for a Burst Descriptor.....	59
Figure 6-17.	Example of UCD Encoded TLV Data.....	62
Figure 6-18.	MAP Format.....	62
Figure 6-19.	MAP Information Element Structure.....	63
Figure 6-20.	Packet PDU Following the Timing Header.....	64
Figure 6-21.	Ranging Response.....	66
Figure 6-22.	Generalized Decision Feedback Equalization Coefficients.....	67
Figure 6-23.	Generalized Equalizer Tap Location Definition.....	68
Figure 6-24.	Example of TLV Data.....	68
Figure 6-25.	Registration Request.....	69
Figure 6-26.	Example of Registration Request Type Value Encodings.....	71
Figure 6-27.	Registration Response Format.....	72
Figure 6-28.	Example of Registration Response Encoding.....	74

Figure 6-29.	Upstream Channel Change Request.....	74
Figure 6-30.	Upstream Channel Change Response	75
Figure 6-31.	Allocation Map	76
Figure 6-32.	Protocol Example.....	80
Figure 7-1.	CM Initialization Overview	88
Figure 7-2.	SDL Notation.....	89
Figure 7-3.	Obtaining Upstream Parameters	91
Figure 7-4.	Message Flows During Scanning and Upstream Parameter Acquisition.....	92
Figure 7-5.	Ranging and Automatic Adjustments Procedure.....	93
Figure 7-6.	Initial Ranging - CM.....	94
Figure 7-7.	Initial Ranging - CM (continued)	95
Figure 7-8.	Initial Ranging - CMTS	96
Figure 7-9.	Establishing IP Connectivity	97
Figure 7-10.	Establishing Time of Day	98
Figure 7-11.	Registration - CMTS.....	100
Figure 7-12.	Registration - CM	101
Figure 7-13.	Periodic Ranging - CMTS	103
Figure 7-14.	Periodic Ranging - CM View	104
Figure 7-15.	Changing Upstream Channels: CMTS View	105
Figure 7-16.	Changing Upstream Channels: CM View.....	106
Figure C-1.	Binary Configuration File Format	117
Figure C-2.	Create TLV Entries for Parameters Required by the CM	118
Figure C-3.	Add CM MIC.....	119
Figure C-4.	Add CMTS MIC	119
Figure C-5.	Add End of Data Marker	119
Figure G-1.	One Downstream and Four Upstream Channels.....	154
Figure H-1.	Spanning Tree Topology	155
Figure H-2.	Spanning Tree Across CMTSes.....	156
Figure N-1.	QPSK symbol mapping	180
Figure N-2.	16QAM Gray-coded symbol mapping	180
Figure N-3.	16QAM differential-coded symbol mapping.....	181
Figure N-4.	Scrambler structure	183
Figure N-5.	Nominal burst timing.....	187
Figure N-6.	Worst-case burst timing.....	188
Figure N-7.	Example frame structures with flexible burst length mode	194
Figure N-8.	Signal-processing sequence	195
Figure N-9.	TDMA upstream transmission processing.....	196
Figure N-10.	Example of interleaving MPEG packets in downstream	201
Figure N-11.	Format of an MPEG Packet.....	202
Figure N-12.	Packet format where a MAC Frame immediately follows the pointer_field	203
Figure N-13.	Packet format with MAC Frame preceded by stuffing bytes.....	203
Figure N-14.	Packet format showing multiple MAC frames in a single packet.....	204
Figure N-15.	Packet format where a MAC Frame spans multiple packets	204

List of Tables

Table 2-1.	Assumed Downstream RF Channel Transmission Characteristics	8
Table 2-2.	Assumed Upstream RF Channel Transmission Characteristics	9
Table 4-1.	I/Q Mapping	19
Table 4-2.	Derivation of Currently Transmitted Symbol Quadrant	20
Table 4-3.	Maximum Channel Width	21
Table 4-4.	Burst Profile Attributes	24
Table 4-5.	User Unique Burst Parameters	24
Table 4-6.	Spurious Emissions	28
Table 4-7.	Adjacent Channel Spurious Emissions	29
Table 4-8.	Spurious Emissions in 5 to 42 MHz	29
Table 4-9.	Maximum Range of Commanded Nominal Receive Power in Each Carrier	33
Table 4-10.	Electrical Output from CM	33
Table 4-11.	Interleaver Characteristics	34
Table 4-12.	CMTS Output	35
Table 4-13.	Electrical Input to CM	36
Table 4-14.	256QAM CM BER Performance	36
Table 5-1.	MPEG Header Format for DOCSIS Data-Over-Cable Packets	39
Table 6-1.	Generic MAC Header Format	46
Table 6-2.	FC Field Format	46
Table 6-3.	Packet PDU Format	48
Table 6-4.	ATM Cell MAC Frame Format	49
Table 6-5.	Reserved PDU Format	50
Table 6-6.	Timing MAC Header Format	51
Table 6-7.	Management MAC Header Format	51
Table 6-8.	Request MAC Header (REQ) Format	52
Table 6-9.	Concatenated MAC Frame Format	53
Table 6-10.	Extended Header Format	54
Table 6-11.	EH Element Format	54
Table 6-12.	EH Element Format	55
Table 6-13.	MAC Management Message Types	57
Table 6-14.	Channel TLV Parameters	59
Table 6-15.	Upstream Physical-Layer Burst Attributes	61
Table 6-16.	Allocation MAP Information Elements (IE)	64
Table 6-17.	Sample Service Class Data Encoding	74
Table 6-18.	Transmit Opportunity	82
Table 6-19.	Example Relating Mini-Slots to Time Ticks	85
Table 7-1.	Recovery Process on Loss of Specific MAC Messages	107
Table 9-1.	Channel Burst Parameters for Advanced Modems	111
Table C-1.	Sample Class of Service Encoding	124
Table C-2.	Sample Capability Encoding	125
Table C-3.	Maximum Number of CPEs	130
Table C-4.	TFTP Server Timestamp	130
Table C-5.	TFTP Server Provisioned Modem Address	130
Table C-6.	Default Configuration Values	130
Table E-1.	Example Channel Burst Parameter Values for QPSK Operation at 160, 320, and 640 ksym/sec	138
Table E-2.	Example Channel Burst Parameter Values for QPSK Operation at 1.28 and 2.56 Msym/sec	139
Table E-3.	Example Channel Burst Parameter Values for 16QAM Operation at 160, 320, and 640 ksym/sec	140
Table E-4.	Example Channel Burst Parameter Values for 16QAM Operation at 1.28 and 2.56 Msym/sec	141
Table E-5.	Example Channel Burst Parameter Values for Power-Up and Acquisition in a New Channel	142
Table E-6.	Frame Format Examples with QPSK Operation	143
Table E-6.	Frame Format Examples with QPSK Operation (continued)	143
Table E-6.	Frame Format Examples with QPSK Operation (continued)	144
Table E-7.	Frame Format Examples with 16QAM Operation	146
Table E-7.	Frame Format Examples with 16QAM Operation (continued)	147

Table E-7.	Frame Format Examples with 16QAM Operation (continued)	147
Table F-1.	Broadcasting Allocations Between 5 and 42 MHz.....	150
Table F-2.	Amateur and Citizens Band Allocations Between 5 and 42 MHz.....	151
Table F-3.	Data Carriers in the Gaps Between Broadcasting Bands.....	152
Table F-4.	Data Carriers in the Gaps Between Broadcasting, Amateur and CB Bands.....	153
Table H-1.	CM Path Cost.....	157
Table I-1.	Error Codes for MAC Management Messages.....	158
Table N-1.	Assumed downstream RF channel transmission characteristics for analog TV and sound signals...	176
Table N-2.	Assumed upstream RF channel transmission characteristics.....	177
Table N-3.	I/Q mapping.....	179
Table N-4.	Derivation of currently-transmitted symbol quadrant.....	181
Table N-5.	Maximum channel width	182
Table N-6.	Burst Profile attributes	185
Table N-7.	User Unique Burst parameters.....	185
Table N-8.	Spurious emissions	190
Table N-9.	Adjacent channel spurious emissions	191
Table N-10.	Spurious emissions in 5 to 65 MHz.....	191
Table N-11.	Maximum range of commanded nominal received power in each carrier.....	196
Table N-12.	Electrical output from CM.....	197
Table N-13.	Interleaver characteristics	197
Table N-14.	CMTS output	198
Table N-15.	Electrical input to CM.....	199
Table N-16.	256QAM CM BER performance.....	200
Table N-17.	MPEG Header format for EuroDOCSIS Data-Over-Cable packets.....	202

This page intentionally left blank.

1 SCOPE AND PURPOSE

1.1 Scope

This document defines the radio-frequency interface specifications for high-speed data-over-cable systems. They were developed by Cable Television Laboratories (CableLabs) for the benefit of the cable industry, including contributions by operators and vendors from North America, Europe, and other regions.

There are differences in the cable spectrum planning practices adopted for different networks in the world. Therefore two options for physical-layer technology are included, which have equal priority and are not required to be interoperable. One technology option is based on the downstream multi-program television distribution that is deployed in North America using 6 MHz channeling, and supports upstream transmission in the region 5-42 MHz. The other technology option is based on the corresponding European multi-program television distribution and supports upstream in the region 5-65 MHz. Although both options have the same status, the first option was documented earlier and the second option introduced at a later time as an amendment, resulting in the document structure not reflecting this equal priority. The first of these options is defined in Sections 2, 4, 5, Appendix G, and Appendix C.1.1.1, whereas the second is defined by replacing the content of those sections with the content of Appendix N. Correspondingly, [ITU-T.J. 83-B], [NCTA] and [SMS] apply only to the first option, and [EN 300 429] only to the second. Compliance with this document requires compliance with one or other of these implementations, not with both. It is not required that equipment built to one option shall inter-operate with equipment built to the other.

These optional physical-layer technologies allow operators some flexibility within any frequency planning, EMC and safety requirements that are mandated for their area of operation. For example, the 6 MHz downstream based option defined by Sections 2, 4, and 5 might be deployable within an 8 MHz channel plan. Compliance with frequency planning and EMC requirements is not covered by this specification and remains the operators responsibility. In this respect, [FCC15], [FCC76] and [EIA-S542] are relevant to North America and [EN 50081-1], [EN50082-1], [EN 50083-2], [EN 50083-7], and [EN 50083-10] are relevant to the European Community.

Any reference in this document to the transmission of television in the forward channel that is not consistent with [EN 300 429] is outside the normative scope as only [EN 300 429] is used for digital multi-program TV distribution by cable in European applications.

Requirements for safety are outside the scope of the present document. Safety standards for European applications are published by the CENELEC.

Note 1: Examples of such CENELEC product safety standards are [EN 60950] and [EN 50083-1].

Note 2: For CENELEC safety categories of interfaces, see [EG 201 212].

1.2 Requirements

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.

Other text is descriptive or explanatory.

1.3 Background

1.3.1 Service Goals

Cable operators are interested in deploying high-speed data communications systems on cable television systems. Cable Television Laboratories, Inc., and its member companies, have prepared a series of interface specifications that permit the early definition, design, development and deployment of data-over-cable systems on a uniform, consistent, open, non-proprietary, multi-vendor interoperable basis.

The intended service will allow transparent bi-directional transfer of Internet Protocol (IP) traffic, between the cable system headend and customer locations, over an all-coaxial or hybrid fiber/coax (HFC) cable network. This is shown in simplified form in Figure 1-1.




Figure 1-1. Transparent IP Traffic Through the Data-Over-Cable System

The transmission path over the cable system is realized at the headend by a Cable Modem Termination System (CMTS), and at each customer location by a Cable Modem (CM). At the headend (or hub), the interface to the data-over-cable system is called the Cable Modem Termination System - Network-Side Interface (CMTS-NSI) and is specified in [DOCSIS3]. At the customer locations, the interface is called the cable-modem-to-customer-premises-equipment interface (CMCI) and is specified in [DOCSIS4]. The intent is for the operators to transparently transfer IP traffic between these interfaces, including but not limited to datagrams, DHCP, ICMP, and IP Group addressing (broadcast and multicast).

1.3.2 Reference Architecture

Note: This architecture illustrates the North American frequency plans only and is not normative for European applications. Refer to section 1 for applicability.

The reference architecture for the data-over-cable services and interfaces is shown in Figure 1-2.




Figure 1-2. Data-Over-Cable Reference Architecture

This page intentionally left blank.

1.3.3 Server Location

This document refers to several servers which are central to the system operation (e.g., provision and security servers).

The message sequence charts used as examples within this document show sample message exchanges in which access to the servers is via the CMTS. It is important to note that access to these servers need not necessarily be via the CMTS, but MAY be via any CM suitably configured. In this case, the scenarios become slightly more complex, as the message flows are as shown in Figure 1-3. Allowing placement of these components to be at locations other than the CMTS allows the system operator the maximum flexibility in server placement and network configuration. Note that the CMTS MUST be able to initialize without access to the servers in this configuration.




Figure 1-3. Server Location Not At CMTS

2. FUNCTIONAL ASSUMPTIONS

This section describes the characteristics of cable television plant to be assumed for the purposes of operation of the data-over-cable system. It is not a description of CMTS or CM parameters. The data-over-cable system is expected to operate satisfactorily in the environment described in this section.

This section applies to the first technology option referred to in Section 1 (“Scope”). For the second option, refer to Appendix N.

Whenever any reference in this section to frequency plans or compatibility with other services conflicts with any legal requirement for the area of operation, the latter shall take precedence. Any reference to NTSC analog signals in 6 MHz channels does not imply that such signals are physically present.

2.1 Broadband Access Network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid fiber/coax (HFC) network. The generic term “cable network” is used here to cover all cases.

A cable network uses a shared-medium, tree-and-branch architecture with analog transmission. The key functional characteristics assumed in this document are the following:

- Two-way transmission
- A maximum optical/electrical spacing between the CMTS and the most distant customer terminal of 100 miles, although typical maximum separation may be 10-15 miles
- A maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 100 miles, although this would typically be limited to 15 miles

2.2 Equipment Assumptions

2.2.1 Frequency Plan

In the downstream direction, the cable system is assumed to have a passband with a lower edge at 50 or 54 MHz and an upper edge which is implementation-dependent but is typically in the range of 300 to 860 MHz. Within that passband, NTSC analog television signals in 6-MHz channels are assumed to be present on the standard, HRC or IRC frequency plans of EIA Interim Standard IS-6 (later became EIA standard EIA-542), as well as other narrowband and wideband digital signals.

In the upstream direction, the cable system may have a subsplit (5-30 MHz) or extended subsplit (5-42 MHz) passband. NTSC analog television signals in 6-MHz channels may be present, as well as other signals.

2.2.2 Compatibility with Other Services

The CM and CMTS MUST coexist with the other services on the cable network. In particular, they MUST operate satisfactorily in the cable spectrum assigned for CMTS-CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals; and they MUST NOT cause harmful interference to any other services that are assigned to the cable network in spectrum outside of that allocated to the CMTS.

2.2.3 Fault Isolation Impact on Other Users

As the data-over-cable system is a shared-media, point-to-multipoint system, fault-isolation procedures MUST take into account the potential harmful impact of faults and fault-isolation procedures on numerous users of the data-over-cable and other services.

2.3 RF Channel Assumptions

The data-over-cable system, configured with at least one set of defined physical-layer parameters (e.g., modulation, forward error correction, symbol rate, etc.) from the range of configuration settings described in this specification, is expected to be capable of operating with a 1500-byte packet loss rate of less than one percent while forwarding at least 100 packets per second on cable networks having characteristics defined in Section 2.3.

2.3.1 Transmission Downstream

The RF channel transmission characteristics of the cable network in the downstream direction assumed for the purposes of minimal operating capability are described in Table 2-1. This assumes nominal analog video carrier level (peak envelope power) in a 6-MHz channel bandwidth. All conditions are present concurrently.

Table 2-1. Assumed Downstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 50 MHz to as high as 860 MHz. However, the values in this table apply only at frequencies \square 88 MHz.
RF channel spacing (design bandwidth)	6 MHz
Transit delay from headend to most distant $\square\square\square\square\square\square\square$	\square 0.800 msec (typically much less)
Carrier-to-noise ratio in a 6-MHz band (analog video level)	Not less than 35 dB (Note 4)
Carrier-to-interference ratio for total power (discrete and broadband ingress signals)	Not less than 35 dB within the design bandwidth
Composite triple beat distortion for analog modulated carriers	Not greater than -50 dBc within the design bandwidth
Composite second order distortion for analog modulated carriers	Not greater than -50 dBc within the design bandwidth
Cross-modulation level	Not greater than -40 dBc within the design bandwidth
Amplitude ripple	0.5 dB within the design bandwidth
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth
Micro-reflections bound for dominant echo	-10 dBc @ \square 0.5 μ sec, -15 dBc @ \square 1.0 μ sec -20 dBc @ \square 1.5 μ sec, -30 dBc @ \square > 1.5 μ sec
Carrier hum modulation	Not greater than -26 dBc (5%)
Burst noise	Not longer than 25 μ sec at a 10 Hz average rate
Seasonal and diurnal signal level variation	8 dB
Signal level slope, 50-750 MHz	16 dB
Maximum analog video carrier level at the CM input, inclusive of above signal level variation	17 dBmV
Lowest analog video carrier level at the CM input, inclusive of above signal level variation	-5 dBmV

Notes to Table 2-1

1. Transmission is from the headend combiner to the CM input at the customer location.
2. For measurements above the normal downstream operating frequency band (except hum), impairments are referenced to the highest-frequency NTSC carrier level.
3. For hum measurements above the normal downstream operating frequency band, a continuous-wave carrier is sent at the test frequency at the same level as the highest-frequency NTSC carrier.
4. This presumes that the digital carrier is operated at analog peak carrier level. When the digital carrier is operated below the analog peak carrier level, this C/N may be less.
5. Measurement methods defined in [NCTA] or [CableLabs2].

2.3.2 Transmission Upstream

The RF channel transmission characteristics of the cable network in the upstream direction assumed for the purposes of minimal operating capability are described in Table 2-2. All conditions are present concurrently.

Table 2-2. Assumed Upstream RF Channel Transmission Characteristics

Parameter	Value
Frequency range	5 to 42 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	□ 0.800 msec (typically much less)
Carrier-to-noise ratio	Not less than 25 dB
Carrier-to-ingress power (the sum of discrete and broadband ingress signals) ratio	Not less than 25 dB (Note 2)
Carrier-to-interference (the sum of noise, distortion, common-path distortion and cross-modulation) ratio	Not less than 25 dB
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 µsec at a 1 kHz average rate for most cases (Notes 3, 4 and 5)
Amplitude ripple	5-42 MHz: 0.5 dB/MHz
Group delay ripple	5-42 MHz: 200 ns/MHz
Micro-reflections -- single echo	-10 dBc @ □ 0.5 µsec -20 dBc @ □ 1.0 µsec -30 dBc @ > 1.0 µsec
Seasonal and diurnal signal level variation	Not greater than 8 dB min to max

Notes to Table 2-2

1. Transmission is from the CM output at the customer location to the headend.
2. Ingress avoidance or tolerance techniques MAY be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 0 dBc [CableLabs1].
3. Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier.
4. CableLabs report containing distribution of return-path burst noise measurements and measurement method is forthcoming.
5. Impulse noise levels more prevalent at lower frequencies (< 15 MHz).

2.3.2.1 Availability

Typical cable network availability is considerably greater than 99%.

2.4 Transmission Levels

The nominal power level of the downstream CMTS 64QAM signal(s) within a 6-MHz channel is targeted to be in the range -10 dBc to -6 dBc relative to analog video carrier level and will normally not exceed analog video carrier level. The nominal power level of the upstream CM signal(s) will be as low as possible to achieve the required margin above noise and interference. Uniform power loading per unit bandwidth is commonly followed in setting upstream signal levels, with specific levels established by the cable network operator to achieve the required carrier-to-noise and carrier-to-interference ratios.

2.5 Frequency Inversion

There will be no frequency inversion in the transmission path in either the downstream or upstream directions, i.e., a positive change in frequency at the input to the cable network will result in a positive change in frequency at the output.

3 COMMUNICATION PROTOCOLS

This section provides a high-level overview of the communication protocols that **MUST** be used in the data-over-cable system. Detailed specifications for the physical media dependent, downstream transmission, and media access control sublayers are provided in Section 4, Section 5, and Section 6, respectively.

3.1 Protocol Stack

The CM and CMTS operate as forwarding agents and also as end-systems (hosts). The protocol stacks used in these modes differ as shown below.

The principal function of the cable modem system is to transmit Internet Protocol (IP) packets transparently between the headend and the subscriber location. Certain management functions also ride on IP, so that the protocol stack on the cable network is as shown in Figure 3-1 (this does not restrict the generality of IP transparency between the headend and the customer). These management functions include, for example, supporting spectrum management functions and the downloading of software.

3.1.1 CM and CMTS as Hosts

CMs and CMTSes will operate as IP and LLC hosts in terms of IEEE Standard 802 [IEEE802] for communication over the cable network. The protocol stack at the CM and CMTS RF interfaces is shown in Figure 3-1.




Figure 3-1. Protocol Stack on the RF Interface




Figure 3-3. Example Condition for Network Loops

3.1.2.2 CMTS Forwarding Rules

At the CMTS, if link-layer forwarding is used, it MUST conform to the following general 802.1d guidelines:

- Link-layer frames between a given pair of end-stations MUST be delivered in order.
- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.

The address-learning and -aging mechanisms used are vendor-dependent.

If network-layer forwarding is used, the CMTS should conform to IETF Router Requirements [RFC-1812] with respect to its CMTS-RFI and CMTS-NSI interfaces.

Conceptually, the CMTS forwards data packets at two abstract interfaces: between the CMTS-RFI and the CMTS-NSI, and between the upstream and downstream channels. The CMTS MAY use any combination of link-layer (bridging) and network-layer (routing) semantics at each of these interfaces. The methods used at the two interfaces need not be the same.

Forwarding between the upstream and downstream channels within a MAC layer differs from traditional LAN forwarding in that:

- A single channel is simplex, and cannot be considered a complete interface for most protocol (e.g., 802.1d spanning tree, Routing Information Protocol per [RFC-1058]) purposes.
- Upstream channels are essentially point-to-point, whereas downstream channels are shared-media.
- As a public network, policy decisions may override full connectivity.

For these reasons, an abstract entity called the MAC Forwarder exists within the CMTS to provide connectivity between stations within a MAC domain (see Section 3.2).

3.1.2.3 CM Forwarding Rules

Data forwarding through the CM is link-layer bridging with the following specific rules.

3.1.2.3.1 CPE MAC Address Acquisition

- The CM MUST acquire Ethernet MAC addresses of connected CPE devices, either from the provisioning process or from learning, until the CM acquires its maximum number of CPE addresses (a device-dependent value). Once the CM acquires its maximum number of CPE addresses, newly-discovered CPE addresses MUST NOT replace previously acquired addresses. The CM MUST support acquisition of at least one CPE address.
- The CM MUST allow configuration of CPE addresses during the provisioning process (up to its maximum number of CPE addresses) to support configurations in which learning is not practical nor desired.
- Addresses provided during the CM provisioning MUST take preference over learned addresses.
- CPE addresses MUST NOT be aged out.
- On a CM reset (e.g., a power cycle), all learned addresses MUST be discarded (they are not retained in non-volatile storage, to allow modification of user MAC addresses or movement of the CM).

3.1.2.3.2 Forwarding

CM forwarding in both directions MUST conform to the following general 802.1d guidelines:

- Link-layer frames between a given pair of end-stations MUST be delivered in order.
- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.

Cable-network-to-Ethernet forwarding MUST conform to the following specific rules:

- Frames addressed to unknown destinations MUST NOT be forwarded from the cable port to the Ethernet port.
- Broadcast frames MUST be forwarded to the Ethernet port, unless they are from source addresses which are provisioned or learned as supported CPE devices, in which case they MUST NOT be forwarded.
- The CM MAY use IGMP messages from the CPE to determine which multicast frames to accept from the HFC interface for forwarding. If IGMP is not supported, the CM MUST accept all multicast frames from the HFC for forwarding. All accepted multicast frames MUST be forwarded to the CMCI ports in accordance with filtering configuration settings specified by the cable operator's operation and business support systems, unless they are from source addresses which are provisioned or learned as supported CPE devices, in which case they MUST NOT be forwarded.

Ethernet-to-cable-network forwarding MUST conform to the following specific rules:

- Frames addressed to unknown destinations MUST be forwarded from the Ethernet port to the cable port.
- Broadcast frames MUST be forwarded to the cable port.
- Multicast frames MUST be forwarded to the cable port in accordance with filtering configuration settings specified by the cable operator's operations and business support systems.
- Frames from source addresses other than those provisioned or learned as supported CPE devices MUST NOT be forwarded.
- If a single-user CM has acquired a MAC address (see Section 3.1.2.3.1), it MUST NOT forward data from a second source. Other (non-supported) CPE source addresses MUST be learned from the Ethernet port and this information used to filter local traffic as in a traditional learning bridge.
- If a single-user CM has acquired MAC address A as its supported CPE device and learned B as a second device connected to the Ethernet port, it MUST filter any traffic from A to B.

3.2 The MAC Forwarder

The MAC Forwarder is a MAC sublayer that resides on the CMTS just below the MAC service access point (MSAP) interface, as shown in Figure 3-4. It is responsible for delivering upstream frames to:

- One or more downstream channels
- The MSAP interface

In Figure 3-4, the LLC sublayer and link security sublayers of the upstream and downstream channels on the cable network terminate at the MAC Forwarder.

The MSAP interface user MAY be the NSI-RFI Forwarding process or the CMTS's host protocol stack.




Figure 3-4. MAC Forwarder

Delivery of frames may be based on data-link-layer (bridging) semantics, network-layer (routing) semantics, or some combination. Higher-layer semantics may also be employed (e.g., filters on UDP port numbers). The CMTS **MUST** provide IP connectivity between hosts attached to cable modems, and must do so in a way that meets the expectations of Ethernet-attached customer equipment. For example, the CMTS must either forward ARP packets or it must facilitate a proxy ARP service. The CMTS MAC Forwarder **MAY** provide service for non-IP protocols.

Note that there is no requirement that all upstream and downstream channels be aggregated under one MSAP as shown above. The vendor could just as well choose to implement multiple MSAPs, each with a single upstream and downstream channel.

3.2.1 Example Rules for Data-Link-Layer Forwarding

If the MAC Forwarder is implemented using only data-link-layer semantics, the requirements in this section apply.

Delivery of frames is dependent on the Destination Address within the frame. The means of learning the location of each address is vendor-dependent, and **MAY** include:

- Transparent-bridging-like source-address learning and aging
- Gleaning from MAC Registration Request messages
- Administrative means

If the destination address of a frame is unicast, and that address is associated with a particular downstream channel, the frame **MUST** be forwarded to that channel.²

If the destination address of a frame is unicast, and that address is known to reside on the other (upper) side of the MSAP interface, the frame **MUST** be delivered to the MSAP interface.

If the destination address is broadcast, multicast³, or unknown, the frame **MUST BE** delivered to both the MSAP and to all downstream channels.

Delivery rules are similar to those for transparent bridging:

- Frames from a specific source to a particular destination **MUST** be delivered in order.
- Frames **MUST NOT** be duplicated.
- Frames that cannot be delivered in a timely fashion **MUST** be discarded.
- The Frame Check Sequence **SHOULD** be preserved rather than regenerated.

² Vendors may implement extensions, similar to static addresses in 802.1d/ISO 10038 bridging, that cause such frames to be filtered or handled in some other manner.

³ Note: all multicasts, including 802.1d/ISO 10038 Spanning Tree Bridge BPDUs, **MUST** be forwarded, except for those addressed to the all-CMTSes multicast address which **MUST NOT** be forwarded.

3.3 Network Layer

As stated above, the purpose of the data-over-cable system is to transport IP traffic transparently through the system.

The Network Layer protocol is the Internet Protocol (IP) version 4, as defined in [RFC-791], and migrating to IP version 6.

This document imposes no requirements for reassembly of IP packets.

3.4 Above the Network Layer

Subscribers will be able to use the transparent IP capability as a bearer for higher-layer services. Use of these services will be transparent to the CM.

In addition to the transport of user data, there are several network management and operation capabilities which depend upon the Network Layer. These include:

- SNMP (Simple Network Management Protocol, [RFC-1157]), for network management
- TFTP (Trivial File Transfer Protocol, [RFC-1350]), a file transfer protocol, for downloading software and configuration information
- DHCP (Dynamic Host Configuration Protocol, DHCP [RFC-2131]), a framework for passing configuration information to hosts on a TCP/IP network

DHCP, TFTP, and ToD client messages generated by the CM MUST only be sent via the RF Interface. DHCP, TFTP and ToD client messages include DHCPDISCOVER, DHCPREQUEST, DHCPDECLINE, DHCPRELEASE, DHCPINFORM, TFTP-RRQ, TFTP-ACK, and ToD request.

The CM's DHCP, TFTP, and ToD client MUST ignore DHCP, TFTP and ToD server message received on the CMCI port. DHCP, TFTP, and ToD server messages included: DHCPPOFFER, DHCPACK, DHCPNAK, TFTP-DATA, and ToD time message.

3.5 Data Link Layer

The Data Link Layer is divided into sublayers in accordance with [IEEE802], with the addition of Link-Layer security in accordance with [DOCSIS8]. The sublayers, from the top, are:

- Logical Link Control (LLC) sublayer (Class 1 only)
- Link-Layer Security sublayer
- Media Access Control (MAC) sublayer

3.5.1 LLC Sublayer

The LLC sublayer MUST be provided in accordance with [ISO/IEC10039]. Address resolution MUST be used as defined in [RFC-826]. The MAC-to-LLC service definition is specified in [ISO/IEC10039].

3.5.2 Link-Layer Security Sublayer

Link-layer security MUST be provided in accordance with [DOCSIS8].

3.5.3 MAC Sublayer

The definition, in detail, of the MAC sublayer and associated interfaces is provided in Section 6 of this document.

The MAC sublayer defines a single transmitter for each downstream channel — the CMTS. All CMs listen to all frames transmitted on the downstream channel upon which they are registered and accept those where the destination matches the CM itself or CPEs reached via the CMCI port. CMs can communicate with other CMs only through the CMTS.

The upstream channel is characterized by many transmitters (CMs) and one receiver (the CMTS). Time in the upstream channel is slotted, providing for Time Division Multiple Access at regulated time ticks. The CMTS provides the time reference and controls the allowed usage for each interval. Intervals may be granted for transmissions by particular CMs, or for contention by all CMs. CMs may contend to request transmission time. To a limited extent, CMs may also contend to transmit actual data. In both cases, collisions can occur and retries are used.

Section 6 describes the MAC-sublayer messages from the CMTS which direct the behavior of the CMs on the upstream channel, as well as messaging from the CMs to the CMTS.

3.5.3.1 Overview

Some of the MAC protocol highlights include:

- Bandwidth allocation controlled by CMTS
- A stream of mini-slots in the upstream
- Dynamic mix of contention- and reservation-based upstream transmit opportunities
- Bandwidth efficiency through support of variable-length packets
- Extensions provided for future support of ATM or other Data PDU
- Class-of-service support
- Extensions provided for security at the Data Link layer
- Support for a wide range of data rates

3.5.3.2 MAC Service Definition

The MAC sublayer service definition is in Appendix D.

3.6 Physical Layer

The Physical (PHY) layer is comprised of two sublayers:

- Transmission Convergence sublayer (present in the downstream direction only)
- Physical Media Dependent (PMD) sublayer

3.6.1 Downstream Transmission Convergence Sublayer

The Downstream Transmission Convergence sublayer exists in the downstream direction only. It provides an opportunity for additional services over the physical-layer bitstream. These additional services might include, for example, digital video. Definition of any such additional services is beyond the scope of this document.

This sublayer is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets, each consisting of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the data-over-cable MAC. Other values of the header may indicate other payloads. The mixture of payloads is arbitrary and controlled by the CMTS.

The Downstream Transmission Convergence sublayer is defined in Section 5 of this document.

3.6.2 PMD Sublayer

3.6.2.1 Overview

The PMD sublayer involves digitally modulated RF carriers on the analog cable network.

In the downstream direction, the PMD sublayer is based on [ITU J.83-B], with the exceptions called out in Section 4.3.2, and includes these features:

- 64 and 256 QAM modulation formats
- 6-MHz occupied spectrum coexists with all other signals on the cable plant
- Concatenation of Reed-Solomon block code and Trellis code supports operation in a higher percentage of North American cable networks
- Variable-depth interleaver supports both latency-sensitive and -insensitive data

The features in the upstream direction are as follows:

- Flexible and programmable CM under control of the CMTS
- Frequency agility
- Time division multiple access
- QPSK and 16 QAM modulation formats
- Support of both fixed-frame and variable-length PDU formats
- Multiple symbol rates
- Programmable Reed-Solomon block coding
- Programmable preambles

3.6.2.2 Interface Points

Three RF interface points are defined at the PMD sublayer:

- a) Downstream output on the CMTS
- b) Upstream input on the CMTS
- c) Cable in/out at the cable modem

Separate downstream output and upstream input interfaces on the CMTS are required for compatibility with typical downstream and upstream signal-combining and -splitting arrangements in headends.

4. PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION

This section applies to the first technology option referred to in section 1 (“scope”). For the second option, refer to Appendix N.

Whenever any reference in this section to spurious emissions conflicts with any legal requirement for the area of operation, the latter shall take precedence.

4.1 Scope

This specification defines the electrical characteristics and protocol for a cable modem (CM) and cable modem termination system (CMTS). It is the intent of this specification to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of this specification to imply any specific implementation.

4.2 Upstream

4.2.1 Overview

The upstream Physical Media Dependent (PMD) sublayer uses a FDMA/TDMA burst modulation format, which provides five symbol rates and two modulation formats (QPSK and 16QAM). The modulation format includes pulse shaping for spectral efficiency, is carrier-frequency agile, and has selectable output power level. The PMD sublayer format includes a variable-length modulated burst with precise timing beginning at boundaries spaced at integer multiples of 6.25 μ sec apart (which is 16 symbols at the highest data rate).

Each burst supports a flexible modulation, symbol rate, preamble, randomization of the payload, and programmable FEC encoding.

All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis.

The PMD sublayer can support a near-continuous mode of transmission, wherein ramp-down of one burst MAY overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. The system timing of the TDMA transmissions from the various CMs MUST provide that the center of the last symbol of one burst and the center of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard time MUST be greater than or equal to the duration of five symbols plus the maximum timing error. Timing error is contributed by both the CM and CMTS. CM timing performance is specified in Section 4. Maximum timing error and guard time may vary with CMTSes from different vendors.

The upstream modulator is part of the cable modem which interfaces with the cable network. The modulator contains the actual electrical-level modulation function and the digital signal-processing function; the latter provides the FEC, preamble prepend, symbol mapping, and other processing steps. This specification is written with the idea of buffering the bursts in the signal processing portion, and with the signal processing portion (1) accepting the information stream a burst at a time, (2) processing this stream into a complete burst of symbols for the modulator, and (3) feeding the properly-timed bursted symbol stream to a memoryless modulator at the exact burst transmit time. The memoryless portion of the modulator only performs pulse shaping and quadrature upconversion.

At the Demodulator, similar to the Modulator, there are two basic functional components: the demodulation function and the signal processing function. Unlike the Modulator, the Demodulator resides in the CMTS and the specification is written with the concept that there will be one demodulation function (not necessarily an actual physical demodulator) for each carrier frequency in use. The demodulation function would receive all bursts on a given frequency.

Note: Unit design approach should be cognizant of the multiple-channel nature of the demodulation and signal processing to be carried out at the headend, and partition or share functionality appropriately to optimally leverage the multi-channel application. A Demodulator design supporting multiple channels in a Demodulator unit may be appropriate.

The demodulation function of the Demodulator accepts a varying-level signal centered around a commanded power level and performs symbol timing and carrier recovery and tracking, burst acquisition, and demodulation. Additionally, the demodulation function provides an estimate of burst timing relative to a reference edge, an estimate of received signal power, an estimate of signal-to-noise ratio, and may engage adaptive equalization to mitigate the effects of a) echoes in the cable plant, b) narrowband ingress and c) group delay. The signal-processing function of the Demodulator performs the inverse processing of the signal-processing function of the Modulator. This includes accepting the demodulated burst data stream and decoding, etc., and possibly multiplexing the data from multiple channels into a single output stream. The signal-processing function also provides the edge-timing reference and gating-enable signal to the demodulators to activate the burst acquisition for each assigned burst slot. The signal-processing function may also provide an indication of successful decoding, decoding error, or fail-to-decode for each codeword and the number of corrected Reed-Solomon symbols in each codeword. For every upstream burst, the CMTS has a prior knowledge of the exact burst length in symbols (see Sections 4.2.6, 4.2.10.1, and A.2).

4.2.2 Modulation Formats

The upstream modulator **MUST** provide both QPSK and 16QAM modulation formats.

The upstream demodulator **MUST** support QPSK, 16QAM, or both modulation formats.

4.2.2.1 Modulation Rates

The upstream modulator **MUST** provide QPSK at 160, 320, 640, 1,280, and 2,560 ksym/sec, and 16QAM at 160, 320, 640, 1,280, and 2,560 ksym/sec.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrowband ingress, as discussed in Appendix F.

The upstream symbol rate **MUST** be fixed for each upstream frequency.

4.2.2.2 Symbol Mapping

The modulation mode (QPSK or 16QAM) is programmable. The symbols transmitted in each mode and the mapping of the input bits to the I and Q constellation **MUST** be as defined in Table 4-1. In the table, I₁ is the MSB of the symbol map, Q₁ is the LSB for QPSK, and Q₀ is the LSB for 16QAM. Q₁ and I₀ have intermediate bit positions in 16QAM. The MSB **MUST** be the first bit in the serial data into the symbol mapper.

Table 4-1. I/Q Mapping

QAM Mode	Input bit Definitions
QPSK	I ₁ Q ₁
16QAM	I ₁ Q ₁ I ₀ Q ₀

The upstream QPSK symbol mapping **MUST** be as shown in Figure 4-1.




Figure 4-1. QPSK Symbol Mapping

The 16QAM non-inverted (Gray-coded) symbol mapping MUST be as shown in Figure 4-2.




Figure 4-2. 16QAM Gray-Coded Symbol Mapping

The 16QAM differential symbol mapping MUST be as shown in Figure 4-3




Figure 4-3. 16QAM Differential-Coded Symbol Mapping

If differential quadrant encoding is enabled, the currently-transmitted symbol quadrant is derived from the previously transmitted symbol quadrant and the current input bits via Table 4-2.

Table 4-2. Derivation of Currently Transmitted Symbol Quadrant

Current Input Bits I(1) Q(1)	Quadrant Phase Change	MSBs of Previously Transmitted Symbol	MSBs for Currently Transmitted Symbol
00	0°	11	11
00	0°	01	01
00	0°	00	00
00	0°	10	10
01	90°	11	01
01	90°	01	00
01	90°	00	10
01	90°	10	11
11	180°	11	00
11	180°	01	10
11	180°	00	11
11	180°	10	01
10	270°	11	10
10	270°	01	11
10	270°	00	01
10	270°	10	00

4.2.2.3 Spectral Shaping

The upstream PMD sublayer MUST support a 25% Nyquist square root raised cosine shaping.

The occupied spectrum MUST NOT exceed the channel widths shown in Table 4-3.

Table 4-3. Maximum Channel Width

Symbol Rate (ksym/sec)	Channel Width (kHz) ^a
160	200
320	400
640	800
1,280	1,600
2,560	3,200

^a Channel width is the -30 dB bandwidth.

4.2.2.4 Upstream Frequency Agility and Range

The upstream PMD sublayer MUST support operation over the frequency range of 5-42 MHz edge to edge.

Offset frequency resolution MUST be supported having a range of ± 32 kHz (increment = 1 Hz; implement within ± 10 Hz).

4.2.2.5 Spectrum Format

The upstream modulator MUST provide operation with the format $s(t) = I(t) \cdot \cos(\omega t) - Q(t) \cdot \sin(\omega t)$, where t denotes time and ω denotes angular frequency.

4.2.3 FEC Encode

4.2.3.1 FEC Encode Modes

The upstream modulator MUST be able to provide the following selections: Reed-Solomon codes over GF(256) with $T = 1$ to 10 or no FEC coding.

The following Reed-Solomon generator polynomial MUST be supported:

$$g(x) = (x + \alpha^0)(x + \alpha^1) \dots (x + \alpha^{2T-1}) \text{ where the primitive element alpha is } 0x02 \text{ hex}$$

The following Reed-Solomon primitive polynomial MUST be supported:

$$p(x) = x^8 + x^4 + x^3 + x^2 + 1$$

The upstream modulator MUST provide codewords from a minimum size of 18 bytes (16 information bytes [k] plus two parity bytes for $T = 1$ error correction) to a maximum size of 255 bytes (k-bytes plus parity-bytes). The uncoded word can have a minimum size of one byte.

In Shortened Last Codeword mode, the CM MUST provide the last codeword of a burst shortened from the assigned length of k data bytes per codeword as described in Section 4.2.10.1.2 of this document.

The value of T MUST be configured in response to the Upstream Channel Descriptor from the CMTS.

4.2.3.2 FEC Bit-to-Symbol Ordering

The input to the Reed-Solomon Encoder is logically a serial bit stream from the MAC layer of the CM, and the first bit of the stream **MUST** be mapped into the MSB of the first Reed-Solomon symbol into the encoder. The MSB of the first symbol out of the encoder **MUST** be mapped into the first bit of the serial bit stream fed to the Scrambler.

[Note that the MAC byte-to-serial upstream convention calls for the byte LSB to be mapped into the first bit of the serial bit stream per Section 6.2.1.3.]

4.2.4 Scrambler (Randomizer)

The upstream modulator **MUST** implement a scrambler (shown in Figure 4-4) where the 15-bit seed value **MUST** be arbitrarily programmable.

At the beginning of each burst, the register is cleared and the seed value is loaded. The seed value **MUST** be used to calculate the scrambler bit which is combined in an XOR with the first bit of data of each burst (which is the MSB of the first symbol following the last symbol of the preamble).

The scrambler seed value **MUST** be configured in response to the Upstream Channel Descriptor from the CMTS.

The polynomial **MUST** be $x^{15} + x^{14} + 1$.




Figure 4-4. Scrambler Structure

4.2.5 Preamble Prepend

The upstream PMD sublayer **MUST** support a variable-length preamble field that is prepended to the data after they have been randomized and Reed-Solomon encoded.

The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is I1 in the first symbol of the burst (see Section 4.2.2.2). The first bit of the Preamble Pattern is designated by the Preamble Value Offset as described in Section 6.3.2.2, Table 6-15.

The value of the preamble that is prepended **MUST** be programmable and the length **MUST** be 0, 2, 4, ..., or 1024 bits for QPSK and 0, 4, 8, ..., or 1024 bits for 16QAM. Thus, the maximum length of the preamble is 512 QPSK symbols or 256 QAM symbols.

The preamble length and value MUST be configured in response to the Upstream Channel Descriptor message transmitted by the CMTS.

4.2.6 Transmit Pre-Equalizer

The requirements in this section of the pre-equalization are optional for the CM. However, if a CM implements transmit equalization, the CM is required to be in conformance to the requirements in this section.

A transmit pre-equalizer of a linear equalizer structure, as shown in Figure 4-5, MUST be configured by the CM in response to the Ranging Response (RNG-RSP) message transmitted by the CMTS. The pre-equalizer MUST support a symbol (T)-spaced equalizer structure with 8 taps. The pre-equalizer MAY have 1, 2, or 4 samples per symbol, with a tap length longer than 8 symbols.




Figure 4-5 Transmit Pre-Equalizer Structure

The RNG-RSP MAC message, (see Section 6.3.2.5) uses 16 bits per coefficient in fractional two's complement notation- "s1.14" (sign bit, integer bit, binary point, and 14 fractional bits) to define the CM transmit equalization information. The CM MUST convolve the coefficients sent by the CMTS with the existing coefficients to get the new coefficients.

In response to an initial ranging request and periodic ranging requests prior to CM registration, when the CMTS sends the pre-equalizer coefficients, the CMTS MUST compute and send them with an equalizer length of 8 and in symbol-spaced format. After registration, the CMTS MAY use a fractionally spaced equalizer format (T/2- or T/4-spaced) with a longer tap length to match the CM pre-equalizer capabilities that the CMTS learned from the REG-REQ message modem capabilities field. See section C.7.7 for proper use of the modem capabilities field.

Prior to making an initial ranging request and whenever the upstream channel frequency or upstream channel symbol rate changes, the CM MUST initialize the coefficients of the pre-equalizer to a default setting in which all coefficients are zero except the real coefficient of the first tap (*i.e.*, F1). During initial ranging, the CM, not the CMTS, MUST compensate for the delay (ranging offset) due to a shift from the first tap to a new main tap location of the equalizer coefficients sent by the CMTS. The pre-equalizer coefficients are then updated through the subsequent ranging process (periodic station maintenance). The CMTS MUST not move the main tap location during periodic station maintenance. Equalizer coefficients may be included in every RNG-RSP message, but typically they only occur when the CMTS determines the channel response has significantly changed. The frequency of equalizer coefficient updates in the RNG-RSP message is determined by the CMTS.

The CM MUST normalize the pre-equalizer coefficients in order to guarantee proper operation (such as not to overflow or clip). The CM MUST also compensate for the change in transmit power due to the gain (or loss) of the new coefficients. If the CM equalizer structure implements the same number of coefficients as assigned in the RNG-RSP message, then the CM MUST not change the location of the main tap in the RNG-RSP message. If the CM equalizer structure implements a different number of coefficients than defined in the RNG-RSP message, the CM

MAY shift the location of the main tap value. Again, in doing so, the CM MUST adjust its ranging offset, in addition to any adjustment in the RNG-RSP message, by an amount that compensates for the movement of the main tap location.

4.2.7 Burst Profiles

The transmission characteristics are separated into three portions: a) Channel Parameters, b) Burst Profile Attributes, and c) User Unique Parameters. The Channel Parameters include i) the symbol rate (five rates from 160 ksym/sec to 2.56 Msym/sec in octave steps), ii) the center frequency (Hz), and iii) the 1024-bit Preamble Superstring. The Channel Parameters are further described in Section 6.3.2.2, Table 6-14; these characteristics are shared by all users on a given channel. The Burst Profile Attributes are listed in Table 4-4, and are further described in Section 6.3.2.2, Table 6-15; these parameters are the shared attributes corresponding to a burst type. The User Unique Parameters may vary for each user even when using the same burst type on the same channel as another user (for example, Power Level), and are listed in Table 4-5.

The CM MUST generate each burst at the appropriate time as conveyed in the mini-slot grants provided by the CMTS MAPs (Section 6.3.2.2).

The CM MUST support all burst profiles commanded by the CMTS via the Burst Descriptors in the UCD (Section 6.3.2.2), and subsequently assigned for transmission in a MAP (Section 6.3.2.3).

Table 4-4. Burst Profile Attributes

Burst Profile Attributes	Configuration Settings
Modulation	QPSK, 16 QAM
Diff Enc	On/Off
Preamble Length	0-1024 bits (Note Section 4.2.5)
Preamble Value offset	0 to 1022
FEC Error Correction (T bytes)	0 to 10 (0 implies FEC = off)
FEC Codeword Information Bytes (k)	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on)
Scrambler Seed	15 bits
Maximum Burst Length (minislots) ^a	0 to 255
Guard Time	5 to 255 symbols
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off

- a. A burst length of 0 mini-slots in the Channel Profile indicates that the burst length is variable on that channel for that burst type. The burst length, while not fixed, is granted explicitly by the CMTS to the CM in the MAP.

Table 4-5. User Unique Burst Parameters

User Unique Parameter	Configuration Settings
Power Level ^a	+8 to +55 dBmV (16QAM) +8 to +58 dBmV (QPSK) 1-dB steps
Offset Frequency ^a	Range = □32 kHz; increment = 1 Hz; implement within □10 Hz
Ranging Offset	0 to (216 - 1), increments of 6.25 μsec/64
Burst Length (mini-slots) if variable on this channel (changes burst-to-burst)	1 to 255 mini-slots

Transmit Equalizer Coefficients(advanced modems only)	Up to 64 coefficients; 4 bytes per coefficient: 2 real and 2 complex
---	--

- a. Values in table apply for this given channel and symbol rate

The CM MUST implement the Offset Frequency to within ± 10 Hz.

Ranging Offset is the delay correction applied by the CM to the CMTS Upstream Frame Time derived at the CM, in order to synchronize the upstream transmissions in the TDMA scheme. The Ranging Offset is an advancement equal to roughly the round-trip delay of the CM from the CMTS. The CMTS MUST provide feedback correction for this offset to the CM, based on reception of one or more successfully received bursts (i.e., satisfactory result from each technique employed: error correction and/or CRC), with accuracy within 1/2 symbol and resolution of 1/64 of the frame tick increment ($6.25 \mu\text{sec}/64 = 0.09765625 \mu\text{sec} = 1/4$ the symbol duration of the highest symbol rate = 10.24 MHz^{-1}). The CMTS sends adjustments to the CM, where a negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM. The CM MUST implement the correction with resolution of at most 1 symbol duration (of the symbol rate in use for a given burst), and (other than a fixed bias) with accuracy within $\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol owing to resolution. The accuracy of CM burst timing of $\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol is relative to the mini-slot boundaries derivable at the CM based on an ideal processing of the timestamp signals received from the CMTS.

The CM MUST be capable of switching burst profiles with no reconfiguration time required between bursts except for changes in the following parameters: 1) Output Power, 2) Modulation, 3) Symbol Rate, 4) Offset frequency, 5) Channel Frequency, and 6) Ranging Offset.

For Symbol Rate, Offset frequency and Ranging Offset, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp-down time of one burst and the ramp-up time of the next burst as well as the overall transmitter delay time including the pipeline delay and optional pre-equalizer delay. For modulation type changes, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol center of one burst and the first symbol center of the following burst.

Output Power, Symbol Rate, Offset frequency, Channel Frequency, and Ranging Offset MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Transmitted Output Power, Symbol Rate, Offset frequency, Channel Frequency and Ranging Offset MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted. The modulation MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted, EXCLUDING the effect of the transmit equalizer (if present in the CM). [This is to be verified with the transmit equalizer providing no filtering; delay only, if that. Note that if the CMTS has decision feedback in its equalizer, it may need to provide more than the 96 symbol gap between bursts of different modulation type which the same CM may use; this is a CMTS decision.] Negative ranging offset adjustments will cause the 96 symbol guard to be violated. The CMTS must assure that this does not happen by allowing extra guard time between bursts that is at least equal to the amount of negative ranging offset.

If Channel Frequency is to be changed, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 100 msec between the last symbol center of one burst and the first symbol of the following burst.

The Channel Frequency of the CM MUST be settled within the phase noise and accuracy requirements of Sections 4.2.9.5 and 4.2.9.6 within 100 msec from the beginning of the change.

If Output Power is to be changed by 1 dB or less, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 5 μsec between the last symbol center of one burst and the first symbol center of the following burst.

If Output Power is to be changed by more than 1 dB, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 10 μsec between the last symbol center of one burst and the first symbol center of the following burst.

The Output Power of the CM MUST be settled to within ± 0.1 dB of its final output power level a) within 5 μ sec from the beginning of a change of 1 dB or less, and b) within 10 μ sec from the beginning of a change greater than 1 dB.

The output transmit power MUST be maintained constant within a TDMA burst to within less than 0.1 dB (excluding the amount theoretically present due to pulse shaping, and amplitude modulation in the case of 16 QAM).

4.2.8 Burst Timing Convention

Figure 4-6 illustrates the nominal burst timing.




Figure 4-6 Nominal Burst Timing

Figure 4-7 indicates worst-case burst timing. In this example, burst N arrives 1.5 symbols late, and burst N+1 arrives 1.5 symbols early, but separation of 5 symbols is maintained; 8-symbol guardband shown.




Figure 4-7. Worst-Case Burst Timing

At a symbol rate of R_s , symbols occur at a rate of one each $T_s = 1/R_s$ seconds. Ramp Up and Ramp Down are the spread of a symbol in the time domain beyond T_s duration owing to the symbol-shaping filter. If only one symbol were transmitted, its duration would be longer than T_s due to the shaping filter impulse response being longer than T_s . The spread of the first and last symbols of a burst transmission effectively extends the duration of the burst to longer than $N * T_s$, where N is the number of symbols in the burst.

4.2.9 Transmit Power Requirements

The upstream PMD sublayer MUST support varying the amount of transmit power. Requirements are presented for 1) the range of commanded transmit power, 2) the step size of the power commands, and 3) the accuracy (actual output power compared to the commanded amount) of the response to the command.

The mechanism by which power adjustments are performed is defined in Section 7.2.4 of this document. Such adjustments MUST be within the ranges of tolerances described below.

4.2.9.1 Output Power Agility and Range

The output transmit power in the design bandwidth MUST be variable over the range of +8 dBmV to 55 dBmV (16 QAM) or 58 dBmV (QPSK), in 1-dB steps. The absolute accuracy of the transmitted power MUST be ± 2 dB, and the step size accuracy ± 0.4 dB, with an allowance for hysteresis while switching in/out a step attenuator (e.g., 20 dB) in which case the accuracy requirement is relaxed to ± 1.4 dB. For example, the actual power increase resulting from a command to increase the power level by 1dB in a CM's next transmitted burst MUST be between 0.6 and 1.4dB.

The step resolution MUST be 1dB or less. When a CM is commanded with finer resolution than it can implement, it MUST round to the nearest supported step size. If the commanded step is half way between two supported step sizes,

the CM MUST choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ± 0.5 dB would result in no step, while a command to step ± 0.75 dB would result in a ± 1 dB step.

4.2.10 Fidelity Requirements

4.2.10.1 Spurious Emissions

The noise and spurious power MUST NOT exceed the levels given in Table 4-6, Table 4-7, and Table 4-8.

In Table 4-6, Inband spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include ISI. The measurement bandwidth for Inband spurious is equal to the symbol rate (e.g., 160 kHz for 160 ksymbols/sec).

The measurement bandwidth for the 3 (or fewer) Carrier-Related Frequency Bands (below 42 MHz) is 160 kHz, with up to 3 160 kHz bands, each with no more than -47 dBc, allowed to be excluded from the “Bands within 5 to 42 MHz Transmitting Burst” specs of Table 4-8.

The measurement bandwidth is also 160 kHz for the Between Bursts specs of Table 4-6 below 42 MHz; the Transmitting Burst specs apply during the mini-slots granted to the CM (when the CM uses all or a portion of the grant), and for a minislot before and after the granted mini-slots. [Note that a minislot may be as short as 32 symbols, or 12.5 microseconds at the 2.56 Msym/sec rate, or as short as 200 microseconds at the 160 ksym/sec rate.] The Between Bursts specs apply except during a used grant of mini-slots, and the minislot before and after the used grant.

Table 4-6. Spurious Emissions

Parameter	Transmitting Burst	Between Bursts
Inband (Inband spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include Inter Symbol Intereference (ISI)).	-40 dBc	The greater of -72 dBc or -59 dBmV
Adjacent Band	See Table 4-7	The greater of -72 dBc or -59 dBmV
3 or Fewer Carrier-Related Frequency Bands(such as second harmonic, if < 42 MHz)	-47 dBc	The greater of -72 dBc or -59 dBmV
Bands within 5 to 42 MHz(excluding assigned channel, adjacent channels, and carrier-related channels)	See Table 4-8	The greater of -72 dBc or -59 dBmV
CM Integrated Spurious Emissions Limits (all in 4 MHz, includes discretes) ^a		
42 to 54 MHz	-max(-40 dBc, -26 dBmV) -	-26 dBmV
54 to 60 MHz	35 dBmV	-40 dBmV
60 to 88 MHz	-40 dBmV	-40 dBmV
88-860 MHz	-45 dBmV	max(-45 dBmV, -40 dBc ^b)
CM Discrete Spurious Emissions Limits ^a		
42 to 54 MHz	-max(-50 dBc, -36 dBmV)	-36 dBmV
54 to 88 MHz	-50 dBmV	-50 dBmV
88 to 860 MHz	-50 dBmV	-50 dBmV

- a. These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST be no greater than -40 dBmV.
- b. “dBc” is relative to the received downstream signal level. Some spurious outputs are proportional to the receive signal level.

4.2.10.1.1 Adjacent Channel Spurious Emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel which could be occupied by a carrier of the same or different symbol rates. The following table lists the required adjacent channel spurious emission levels for all combinations of transmitted carrier symbol rates and adjacent channel symbol rates. The measurement is

performed in an adjacent channel interval that is of appropriate bandwidth and distance from the transmitted carrier based on the symbol rates of the transmitted carrier and the carrier in the adjacent channel.

Table 4-7. Adjacent Channel Spurious Emissions

Transmitted carrier symbol rate	Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier symbol rate
160 Ksym/sec	-45 dBc	20 KHz to 180 KHz	160 Ksym/sec
	-45 dBc	40 KHz to 360 KHz	320 Ksym/sec
	-45 dBc	80 KHz to 720 KHz	640 Ksym/sec
	-42 dBc	160 KHz to 1440 KHz	1280 Ksym/sec
	-39 dBc	320 KHz to 2880 KHz	2560 Ksym/sec
All other symbol rates	-45 dBc	20 KHz to 180 KHz	160 Ksym/sec
	-45 dBc	40 KHz to 360 KHz	320 Ksym/sec
	-45 dBc	80 KHz to 720 KHz	640 Ksym/sec
	-44 dBc	160 KHz to 1440 KHz	1280 Ksym/sec
	-41 dBc	320 KHz to 2880 KHz	sym/sec

4.2.10.1.2 Spurious Emissions in 5 to 42 MHz

Spurious emissions, other than those in an adjacent channel or carrier related emissions listed above, may occur in intervals that could be occupied by other carriers of the same or different symbol rates. To accommodate these different symbol rates and associated bandwidths, the spurious emissions are measured in an interval equal to the bandwidth corresponding to the symbol rate of the carrier that could be transmitted in that interval. This interval is independent of the current transmitted symbol rate.

The following table lists the possible symbol rates that could be transmitted in an interval, the required spurious level in that interval, and the initial measurement interval at which to start measuring the spurious emissions. Measurements should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge, 5 MHz or 42 MHz, is reached. Measurement intervals should not include carrier related emissions.

Table 4-8. Spurious Emissions in 5 to 42 MHz

Possible symbol rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 Ksym/sec	-53 dBc	220 KHz to 380 KHz
320 Ksym/sec	-50 dBc	240 KHz to 560 KHz
640 Ksym/sec	-47 dBc	280 KHz to 920 KHz
1280 Ksym/sec	-44 dBc	360 KHz to 1640 KHz
2560 Ksym/sec	-41 dBc	520 KHz to 3080 KHz

4.2.10.2 Spurious Emissions During Burst On/Off Transients

Each transmitter MUST control spurious emissions, prior to and during ramp up and during and following ramp down, before and after a burst in the TDMA scheme.

On/off spurious emissions, such as the change in voltage at the upstream transmitter output due to enabling or disabling transmission, MUST be no more than 100 mV, and such a step MUST be dissipated no faster than 2 μs of constant slewing. This requirement applies when the CM is transmitting at +55 dBmV or more; at backed-off transmit

levels, the maximum change in voltage MUST decrease by a factor of 2 for each 6-dB decrease of power level from +55 dBmV, down to a maximum change of 7 mV at 31 dBmV and below. This requirement does not apply to CM power-on and power-off transients.

4.2.10.3 Symbol Error Rate (SER)

Modulator performance MUST be within 0.5 dB of theoretical SER vs C/N (i.e., E_s/N_0), for SER as low as 10^{-6} uncoded, for QPSK and 16 QAM.

The SER degradation is determined by the cluster variance caused by the transmit waveform at the output of an ideal square-root raised-cosine receive filter. It includes the effects of ISI, spurious, phase noise, and all other transmitter degradations.

Cluster SNR should be measured on a modulation analyzer using a square-root raised cosine receive filter with $\alpha = 0.25$. The measured SNR MUST be better than 30 dB.

4.2.10.4 Filter Distortion

The following requirements assume that any pre-equalization is disabled.

4.2.10.4.1 Amplitude

The spectral mask MUST be the ideal square root raised cosine spectrum with $\alpha = 0.25$, within the ranges given below:

- $f_c - R_s/4$ Hz to $f_c + R_s/4$ Hz: -0.3 dB to +0.3 dB
- $f_c - 3R_s/8$ Hz to $f_c - R_s/4$ Hz, and $f_c + R_s/4$ Hz to $f_c + 3R_s/8$ Hz: -0.5 dB to 0.3 dB
- $f_c - R_s/2$ Hz and $f_c + R_s/2$ Hz: -3.5 dB to -2.5 dB
- $f_c - 5R_s/8$ Hz and $f_c + 5R_s/8$ Hz: no greater than -30 dB

where f_c is the center frequency, R_s is the symbol rate, and the spectral density is measured with a resolution bandwidth of 10 KHz or less.

4.2.10.4.2 Phase

$f_c - 5R_s/8$ Hz to $f_c + 5R_s/8$ Hz: Group Delay Variation MUST NOT be greater than 100 nsec.

4.2.10.5 Carrier Phase Noise

The upstream transmitter total integrated phase noise (including discrete spurious noise) MUST be less than or equal to -43 dBc summed over the spectral regions spanning 1 kHz to 1.6 MHz above and below the carrier.

4.2.10.6 Channel Frequency Accuracy

The CM MUST implement the assigned channel frequency within \square 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

4.2.10.7 Symbol Rate Accuracy

The upstream modulator MUST provide an absolute accuracy of symbol rates \square 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

4.2.10.8 Symbol Timing Jitter

Peak-to-peak symbol jitter, referenced to the previous symbol zero-crossing, of the transmitted waveform, MUST be less than 0.02 of the nominal symbol duration over a 2-sec period. In other words, the difference between the maximum and the minimum symbol duration during the 2-sec period shall be less than 0.02 of the nominal symbol duration for each of the five upstream symbol rates.

The peak-to-peak cumulative phase error, referenced to the first symbol time and with any fixed symbol frequency offset factored out, MUST be less than 0.04 of the nominal symbol duration over a 0.1-sec period. In other words, the difference between the maximum and the minimum cumulative phase error during the 0.1-sec period shall be less than 0.04 of the nominal symbol duration for each of the five upstream symbol rates. Factoring out a fixed symbol frequency offset is to be done by using the computed mean symbol duration during the 0.1 sec.

4.2.11 Frame Structure

Figure 4-8 shows two examples of the frame structure: one where the packet length equals the number of information bytes in a codeword, and another where the packet length is longer than the number of information bytes in one codeword, but less than in two codewords. Example 1 illustrates the fixed codeword-length mode, and Example 2 illustrates the shortened last codeword mode. These modes are defined in Section 4.2.10.1.




Figure 4-8. Example Frame Structures with Flexible Burst Length Mode

4.2.11.1 Codeword Length

When FEC is enabled, the CM operates in either fixed-length codeword mode or in shortened-last codeword mode. The minimum number of information bytes in a codeword in either mode is 16. Shortened-last codeword mode only provides a benefit when the number of bytes in a codeword is greater than the minimum of 16 bytes.

The following descriptions apply to an allocated grant of mini-slots in both contention and non-contention regions. (Allocation of mini-slots is discussed in Section 6 of this document.) The intent of the description is to define rules and conventions such that CMs request the proper number of mini-slots and the CMTS PHY knows what to expect regarding the FEC framing in both fixed codeword length and shortened last codeword modes.

4.2.11.1.1 Fixed Codeword Length

With the fixed-length codewords, after all the data are encoded, zero-fill will occur in this codeword if necessary to reach the assigned k data bytes per codeword, and zero-fill MUST continue up to the point when no additional fixed-

length codewords can be inserted before the end of the last allocated mini-slot in the grant, accounting for FEC parity and guard-time symbols.

4.2.11.1.2 Shortened Last Codeword

As shown in Figure 4-7, let k' = the number of information bytes that remain after partitioning the information bytes of the burst into full-length (k burst data bytes) codewords. The value of k' is less than k . Given operation in a shortened last codeword mode, let k'' = the number of burst data bytes plus zero-fill bytes in the shortened last codeword. In shortened codeword mode, the CM will encode the data bytes of the burst (including MAC Header) using the assigned codeword size (k information bytes per codeword) until 1) all the data are encoded, or 2) a remainder of data bytes is left over which is less than k . Shortened last codewords shall not have less than 16 information bytes, and this is to be considered when CMs make requests of mini-slots. In shortened last codeword mode, the CM will zero-fill data if necessary until the end of the mini-slot allocation, which in most cases will be the next mini-slot boundary, accounting for FEC parity and guard-time symbols. In many cases, only $k'' - k'$ zero-fill bytes are necessary to fill out a mini-slot allocation with $16 \leq k'' \leq k$ and $k' \leq k''$. However, note the following.

More generally, the CM is required to zero-fill data until the point when no additional fixed-length codewords can be inserted before the end of the last allocated mini-slot in the grant (accounting for FEC parity and guard-time symbols), and then, if possible, a shortened last codeword of zero-fill shall be inserted to fit into the mini-slot allocation.

If, after zero-fill of additional codewords with k information bytes, there are less than 16 bytes remaining in the allocated grant of mini-slots, accounting for parity and guard-time symbols, the CM shall not create this last shortened codeword.

4.2.12 Signal Processing Requirements

The signal processing order for each burst packet type MUST be compatible with the sequence shown in Figure 4-9 and MUST follow the order of steps in Figure 4-10.





Figure 4-9. Signal-Processing Sequence






Figure 4-10. TDMA Upstream Transmission Processing

4.2.13 Upstream Demodulator Input Power Characteristics

The maximum total input power to the upstream demodulator MUST NOT exceed 35 dBmV in the 5-42 MHz frequency range of operation.

The intended received power in each carrier MUST be within the values shown in Table 4-9.

Table 4-9. Maximum Range of Commanded Nominal Receive Power in Each Carrier

Symbol Rate (ksym/sec)	Maximum Range (dBmV)
160	-16 to +14
320	-13 to +17
640	-10 to +20
1,280	-7 to +23
2,560	-4 to +26

The demodulator MUST operate within its defined performance specifications with received bursts within ±6 dB of the nominal commanded received power.

4.2.14 Upstream Electrical Output from the CM

The CM MUST output an RF modulated signal with the characteristics delineated in Table 4-10.

Table 4-10. Electrical Output from CM

Parameter	Value
Frequency	5 to 42 MHz edge to edge
Level range (one channel)	+8 to +55 dBmV (16QAM)

	+8 to +58 dBmV (QPSK)
Modulation Type	QPSK and 16QAM
Symbol Rate (nominal)	160, 320, 640, 1,280 and 2,560 ksym/sec
Bandwidth	200, 400, 800, 1,600 and 3,200 kHz
Output impedance	75 ohms
Output Return Loss	> 6 dB (5-42 MHz)
Connector	[SCTE 02 1997] or [IEC 60169-24]

4.3 Downstream

4.3.1 Downstream Protocol

The downstream PMD sublayer MUST conform to ITU-T Recommendations J.83, Annex B for Low-Delay Video Applications [ITU J.83-B], with the exceptions called out in Section 4.3.2.

Note: Any reference in this document to the transmission of television in the forward channel that is not consistent with [EN 300 429] is outside the normative scope as only [EN 300 429] is used for digital multi-program TV distribution by cable in European applications. See sec. 1 ("Scope").

4.3.2 Scalable Interleaving to Support Low Latency

The downstream PMD sublayer MUST support a variable-depth interleaver with the characteristics defined in Table 4-11. The table contains a subset of interleaver modes found in [ITU J.83-B].

Table 4-11. Interleaver Characteristics

I (Number of Taps)	J (Increment)	Burst Protection 64QAM/256QAM	Latency 64QAM/256QAM
8	16	5.9 μ sec/4.1 μ sec	0.22 msec/0.15 msec
16	8	12 μ sec/8.2 μ sec	0.48 msec/0.33 msec
32	4	24 μ sec/16 μ sec	0.98 msec/0.68 msec
64	2	47 μ sec/33 μ sec	2.0 msec/1.4 msec
128	1	95 μ sec/66 μ sec	4.0 msec/2.8 msec

The interleaver depth, which is coded in a 4-bit control word contained in the FEC frame synchronization trailer, always reflects the interleaving in the immediately-following frame. In addition, errors are allowed while the interleaver memory is flushed after a change in interleaving is indicated.

Refer to [ITU J.83-B] for the control bit specifications required to specify which interleaving mode is used.

4.3.3 Downstream Frequency Plan

The downstream frequency plan should comply with Harmonic Related Carrier (HRC), Incremental Related Carrier (IRC) or Standard (STD) North American frequency plans per [IS-6]. However, operation below a center frequency of 91 MHz is not required.

4.3.4 CMTS Output Electrical

The CMTS MUST output an RF modulated signal with the following characteristics defined in Table 4-12.

Table 4-12. CMTS Output

Parameter	Value
Center Frequency (f_c)	91 to 857 MHz \pm 30 kHz ^a
Level	Adjustable over the range 50 to 61 dBmV
Modulation Type	64QAM and 256QAM
Symbol Rate (nominal)	
64QAM	5.056941 Msym/sec
256QAM	5.360537 Msym/sec
Nominal Channel Spacing	6 MHz
Frequency response	
64QAM	~18% Square Root Raised Cosine shaping
256QAM	~12% Square Root Raised Cosine shaping
Total Discrete Spurious Inband ($f_c \pm 3$ MHz)	< -57dBc
Inband Spurious and Noise ($f_c \pm 3$ MHz)	< -48dBc; where channel spurious and noise includes all discrete spurious, noise, carrier leakage, clock lines, synthesizer products, and other undesired transmitter products. Noise within +/- 50kHz of the carrier is excluded.
Adjacent channel ($f_c \pm 3.0$ MHz) to ($f_c \pm 3.75$ MHz)	< -58 dBc in 750 kHz
Adjacent channel ($f_c \pm 3.75$ MHz) to ($f_c \pm 9$ MHz)	< -62 dBc, in 5.25 MHz, excluding up to 3 spurs, each of which must be <-60 dBc when measured in a 10 kHz band
Next adjacent channel ($f_c \pm 9$ MHz) to ($f_c \pm 15$ MHz)	Less than the greater of -65 dBc or -12dBmV in 6MHz, excluding up to three discrete spurs. The total power in the spurs must be < -60dBc when each is measured with 10 kHz bandwidth.
Other channels (47 MHz to 1,000 MHz)	< -12dBmV in each 6 MHz channel, excluding up to three discrete spurs. The total power in the spurs must be < -60dBc when each is measured with 10kHz bandwidth.
Phase Noise	1 kHz - 10 kHz: -33dBc double-sided noise power 10 kHz - 50 kHz: -51dBc double-sided noise power 50 kHz - 3 MHz: -51dBc double-sided noise power
Output Impedance	75 ohms
Output Return Loss	> 14 dB within an output channel up to 750 MHz; > 13 dB in an output channel above 750 MHz
Connector	[SCTE 02 1997] or [IEC 60169-24]

^a \pm 30 kHz includes an allowance of 25 kHz for the largest FCC frequency offset normally built into upconverters

4.3.5 Downstream Electrical Input to CM

The CM MUST accept an RF modulated signal with the following characteristics (Table 4-13.)

Table 4-13. Electrical Input to CM

Parameter	Value
Center Frequency	91 to 857 MHz \pm 30 kHz
Level Range (one channel)	-15 dBmV to +15 dBmV
Modulation Type	64QAM and 256QAM
Symbol Rate (nominal)	5.056941 Msym/sec (64QAM) and 5.360537 Msym/sec (256QAM)
Bandwidth	6 MHz (18% Square Root Raised Cosine shaping for 64QAM and 12% Square Root Raised Cosine shaping for 256QAM)
Total Input Power (40-900 MHz)	<30 dBmV
Input (load) Impedance	75 ohms
Input Return Loss	> 6 dB (88-860 MHz)
Connector	[SCTE 02 1997] or [IEC 60169-24]

4.3.6 CM BER Performance

The bit-error-rate performance of a CM MUST be as described in this section. The requirements apply to the I = 128, J = 1 mode of interleaving.

4.3.6.1 64QAM

4.3.6.1.1 64QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) of 23.5 dB or greater.

4.3.6.1.2 64QAM Image Rejection Performance

Performance as described in Section 4.3.6.1.1 MUST be met with analog or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

4.3.6.1.3 64QAM Adjacent Channel Performance

Performance as described in Section 4.3.6.1.1 MUST be met with digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 4.3.6.1.1 MUST be met with analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 4.3.6.1.1, with an additional 0.2-dB allowance, MUST be met with digital signal at +10 dBc in the adjacent channels.

4.3.6.2 256QAM

4.3.6.2.1 256QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) as shown below.

Table 4-14. 256QAM CM BER Performance

Input Receive Signal Level	E_s/N_o
-6 dBmV to +15 dBmV	30 dB or greater
Less than -6 dBmV down to -15 dBmV	33 dB or greater

4.3.6.2.2 256QAM Image Rejection Performance

Performance as described in Section 4.3.6.2.1 MUST be met with analog or digital signal at +10 dBc in any portion of the RF band other than the adjacent channels.

4.3.6.2.3 256QAM Adjacent Channel Performance

Performance as described in Section 4.3.6.2.1 MUST be met with analog or digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 4.3.6.2.1, with an additional 0.5-dB allowance, MUST be met with analog signal at +10 dBc in the adjacent channels.

Performance as described in Section 4.3.6.2.1, with an additional 1.0-dB allowance, MUST be met with digital signal at +10 dBc in the adjacent channels.

4.3.7 CMTS Timestamp Jitter

The CMTS timestamp jitter must be less than 500 ns peak-to-peak at the output of the Downstream Transmission Convergence Sublayer. This jitter is relative to an ideal Downstream Transmission Convergence Sublayer that transfers the MPEG packet data to the Downstream Physical Media Dependent Sublayer with a perfectly continuous and smooth clock at the MPEG packet data rate. Downstream Physical Media Dependent Sublayer processing MUST NOT be considered in timestamp generation and transfer to the Downstream Physical Media Dependent Sublayer.

Thus, any two timestamps N1 and N2 (N2 > N1) which were transferred to the Downstream Physical Media Dependent Sublayer at times T1 and T2 respectively must satisfy the following relationship:

$$| (N2-N1)f_{CMTS} - (T2-T1) | < 500 \times 10^{-9}$$

In the equation, the value of (N2-N1) is assumed to account for the effect of rollover of the timebase counter, and T1 and T2 represent time in seconds. f_{CMTS} is the actual frequency of the CMTS master timebase and may include a fixed frequency offset from the nominal frequency of 10.24 MHz. This frequency offset is bounded by a requirement further below in this section.

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500ns allocated for jitter at the Downstream Transmission Convergence Sublayer output must be reduced by any jitter that is introduced by the Downstream Physical Media Dependent Sublayer.

The CM is expected to meet the burst timing accuracy requirements in Section 4.2.6 when the time stamps contain this worst-case jitter.

Note: Jitter is the error (i.e., measured) relative to the CMTS Master Clock. (The CMTS Master Clock is the 10.24 MHz clock used for generating the timestamps.)

The CMTS 10.24 MHz Master Clock MUST have frequency stability of ± 5 ppm, drift rate $\pm 10^{-8}$ per second, and edge jitter of ± 10 nsec peak-to-peak (± 5 nsec). [The drift rate and jitter requirements on the CMTS Master Clock implies that the duration of two adjacent segments of 10,240,000 cycles will be within 30 nsec, due to 10 nsec jitter on each segment’s duration, and 10 nsec due to frequency drift. Durations of other counter lengths also may be deduced: adjacent 1,024,000 segments, ± 21 nsec; 1,024,000 length segments separated by 1 10,240,000 cycles, ± 30 nsec; adjacent 102,400,000 segments, ± 120 nsec. The CMTS Master Clock MUST meet such test limits in 99% or more measurements.

5. DOWNSTREAM TRANSMISSION CONVERGENCE SUBLAYER

This section applies to the first technology option referred to in Section 1 (“Scope”). For the second option, refer to Appendix N.

5.1 Introduction

In order to improve demodulation robustness, facilitate common receiving hardware for both video and data, and provide an opportunity for the possible future multiplexing of video and data over the PMD sublayer bitstream defined in Section 4, a sublayer is interposed between the downstream PMD sublayer and the Data-Over-Cable MAC sublayer.

The downstream bitstream is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets. These packets consist of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the Data-Over-Cable MAC. Other values of the header may indicate other payloads. The mixture of MAC payloads and those of other services is optional and is controlled by the CMTS.

Figure 5-1 illustrates the interleaving of Data-Over-Cable (DOC) MAC bytes with other digital information (digital video in the example shown).

header=DOC	DOC MAC payload
header=video	digital video payload
header=video	digital video payload
header=DOC	DOC MAC payload
header=video	digital video payload
header=DOC	DOC MAC payload
header=video	digital video payload
header=video	digital video payload
header=video	digital video payload

Figure 5-1. Example of Interleaving MPEG Packets in Downstream

5.2 MPEG Packet Format

The format of an MPEG Packet carrying DOCSIS data is shown in Figure 5-2. The packet consists of a 4-byte MPEG Header, a pointer_field (not present in all packets) and the DOCSIS Payload.

MPEG Header (4 bytes)	pointer_field (1 byte)	DOCSIS Payload (183 or 184 bytes)
--------------------------	---------------------------	--------------------------------------

Figure 5-2. Format of an MPEG Packet

5.3 MPEG Header for DOCSIS Data-Over-Cable

The format of the MPEG Transport Stream header is defined in Section 2.4 of [ITU-T H.222.0]. The particular field values that distinguish Data-Over-Cable MAC streams are defined in Table 5-1. Field names are from the ITU specification.

The MPEG Header consists of 4 bytes that begin the 188-byte MPEG Packet. The format of the header for use on a DOCSIS Data-Over-Cable PID is restricted to that shown in Table 5-1. The header format conforms to the MPEG standard, but its use is restricted in this specification to NOT ALLOW inclusion of an `adaptation_field` in the MPEG packets.

Table 5-1. MPEG Header Format for DOCSIS Data-Over-Cable Packets

Field	Length (bits)	Description
<code>sync_byte</code>	8	0x47; MPEG Packet Sync byte
<code>transport_error_indicator</code>	1	Indicates an error has occurred in the reception of the packet. This bit is reset to zero by the sender, and set to one whenever an error occurs in transmission of the packet
<code>payload_unit_start_indicator</code>	1	A value of one indicates the presence of a <code>pointer_field</code> as the first byte of the payload (fifth byte of the packet)
<code>transport_priority</code>	1	Reserved; set to zero
PID (see Note)	13	DOCSIS Data-Over-Cable well-known PID (0x1FFE)
<code>transport_scrambling_control</code>	2	Reserved, set to '00'
<code>adaptation_field_control</code>	2	'01'; use of the <code>adaptation_field</code> is NOT ALLOWED on the DOCSIS PID
<code>continuity_counter</code>	4	Cyclic counter within this PID

Note: In the future, additional PIDs MAY be assigned to a CM. See Section 9.3 of this document.

5.4 MPEG Payload for DOCSIS Data-Over-Cable

The MPEG payload portion of the MPEG packet will carry the DOCSIS MAC frames. The first byte of the MPEG payload will be a 'pointer_field' if the `payload_unit_start_indicator` (PUSI) of the MPEG header is set.

stuff_byte

This standard defines a `stuff_byte` pattern having a value (0xFF) that is used within the DOCSIS payload to fill any gaps between the DOCSIS MAC frames. This value is chosen as an unused value for the first byte of the DOCSIS MAC frame. The 'FC' byte of the MAC Header will be defined to never contain this value. (FC_TYPE = '11' indicates a MAC-specific frame, and FC_PARM = '11111' is not currently used and, according to this specification, is defined as an illegal value for FC_PARM.)

pointer_field

The `pointer_field` is present as the fifth byte of the MPEG packet (first byte following the MPEG header) whenever the PUSI is set to one in the MPEG header. The interpretation of the `pointer_field` is as follows:

The `pointer_field` contains the number of bytes in this packet that immediately follow the `pointer_field` that the CM decoder must skip past before looking for the beginning of an DOCSIS MAC Frame. A pointer field MUST be present if it is possible to begin a Data-Over-Cable MAC Frame in the packet, and MUST point to either:

1. the beginning of the first MAC frame to start in the packet or
2. to any `stuff_byte` preceding the MAC frame.

5.5 Interaction with the MAC Sublayer

MAC frames may begin anywhere within an MPEG packet, MAC frames may span MPEG packets, and several MAC frames may exist within an MPEG packet.

The following figures show the format of the MPEG packets that carry DOCSIS MAC frames. In all cases, the PUSI flag indicates the presence of the pointer_field as the first byte of the MPEG payload.

Figure 5-3 shows a MAC frame that is positioned immediately after the pointer_field byte. In this case, pointer_field is zero, and the DOCSIS decoder will begin searching for a valid FC byte at the byte immediately following the pointer_field.




Figure 5-3. Packet Format Where a MAC Frame Immediately Follows the pointer_field

Figure 5-4 shows the more general case where a MAC Frame is preceded by the tail of a previous MAC Frame and a sequence of stuffing bytes. In this case, the pointer_field still identifies the first byte after the tail of Frame #1 (a stuff_byte) as the position where the decoder should begin searching for a legal MAC sublayer FC value. This format allows the multiplexing operation in the CMTS to immediately insert a MAC frame that is available for transmission if that frame arrives after the MPEG header and pointer_field have been transmitted.

In order to facilitate multiplexing of the MPEG packet stream carrying DOCSIS data with other MPEG-encoded data, the CMTS SHOULD NOT transmit MPEG packets with the DOCSIS PID which contain only stuff_bytes in the payload area. MPEG null packets SHOULD be transmitted instead. Note that there are timing relationships implicit in the DOCSIS MAC sublayer which must also be preserved by any MPEG multiplexing operation.




Figure 5-4. Packet Format with MAC Frame Preceded by Stuffing Bytes

Figure 5-5 shows that multiple MAC frames may be contained within the MPEG packet. The MAC frames may be concatenated one after the other or be separated by an optional sequence of stuffing bytes.




Figure 5-5. Packet Format Showing Multiple MAC Frames in a Single Packet

Figure 5-6 shows the case where a MAC frame spans multiple MPEG packets. In this case, the pointer_field of the succeeding frame points to the byte following the last byte of the tail of the first frame.




Figure 5-6. Packet Format Where a MAC Frame Spans Multiple Packets

The Transmission Convergence sublayer must operate closely with the MAC sublayer in providing an accurate timestamp to be inserted into the Time Synchronization message (refer to Section 6.3.2.1 and Section 6.5).

5.6 Interaction with the Physical Layer

The MPEG-2 packet stream **MUST** be encoded according to [ITU-T J.83-B], including MPEG-2 transport framing using a parity checksum as described in [ITU-T J.83-B].

5.7 MPEG Header Synchronization and Recovery

The MPEG-2 packet stream **SHOULD** be declared “in frame” (i.e., correct packet alignment has been achieved) when five consecutive correct parity checksums, each 188 bytes from the previous one, have been received.

The MPEG-2 packet stream **SHOULD** be declared “out of frame,” and a search for correct packet alignment started, when nine consecutive incorrect parity checksums are received.

The format of MAC frames is described in detail in Section 6.

6 MEDIA ACCESS CONTROL SPECIFICATION

6.1 Introduction

6.1.1 Overview

This section describes version 1.0 of the DOCSIS MAC protocol. Some of the MAC protocol highlights include:

- Bandwidth allocation controlled by CMTS
- A stream of mini-slots in the upstream
- Dynamic mix of contention-and reservation-based upstream transmit opportunities
- Bandwidth efficiency through support of variable-length packets
- Extensions provided for future support of ATM or other Data PDU
- Class of service support
- Extensions provided for security at the Data Link layer.
- Support for a wide range of data rates.

6.1.2 Definitions

6.1.2.1 MAC-Sublayer Domain

A MAC-sublayer domain is a collection of upstream and downstream channels for which a single MAC Allocation and Management protocol operates. Its attachments include one CMTS and some number of CMs. The CMTS **MUST** service all of the upstream and downstream channels; each CM **MAY** access one or more upstream and downstream channels.

The CMTS **MUST** police and discard any packets received that have a source MAC address that is not a unicast MAC address.

6.1.2.2 MAC Service Access Point

A MAC Service Access Point (MSAP) is an attachment to a MAC-sublayer domain.

6.1.2.3 Service ID

The concept of Service IDs is central to the operation of the MAC protocol. Service IDs provide both device identification and class-of-service management. In particular, they are integral to upstream bandwidth allocation.

A Service ID defines a particular mapping between a CM and the CMTS. This mapping is the basis on which bandwidth is allocated to the CM by the CMTS and by which class of service is implemented. Within a MAC-sublayer domain, all Service IDs **MUST** be unique.

The CMTS **MAY** assign one or more Service IDs (SIDs) to each CM, corresponding to the classes of service required by the CM. This mapping **MUST** be negotiated between the CMTS and the CM during CM registration.

In a basic CM implementation, a single Service ID can be used; for example to offer best-effort IP service. However, the Service ID concept allows for more complex CMs to be developed with support for multiple service classes while supporting interoperability with more basic modems. In particular, the Service ID concept is expected to support the concept of “data flows” on which protocols such as RSVP and RTP are based.

The Service ID is unique within a single MAC-sublayer domain. The length of the Service ID is 14 bits (although the Service ID is sometimes carried in the low-order bits of a 16-bit field).

6.1.2.4 Upstream Intervals, Mini-Slots and 6.25-Microsecond Increments

The upstream transmission time-line is divided into intervals by the upstream bandwidth allocation mechanism. Each interval is an integral number of mini-slots. A “mini-slot” is the unit of granularity for upstream transmission opportunities. There is no implication that any PDU can actually be transmitted in a single mini-slot. Each interval is labeled with a usage code which defines both the type of traffic that can be transmitted during that interval and the physical-layer modulation encoding. A mini-slot is a power-of-two multiple of 6.25- μ s increments, i.e., 2, 4, 8, 16, 32, 64, or 128 times 6.25 μ s. The relationship between mini-slots, bytes, and time ticks is described further in Section 6.5.4. The usage code values are defined in Table 6-16 and allowed use is defined in Section 6.3. The binding of these values to physical-layer parameters is defined in Table 6-14.

6.1.2.5 Frame

A frame is a unit of data exchanged between two (or more) entities at the Data Link Layer. A MAC frame consists of a MAC Header (beginning with a Frame Control byte; see Figure 6-3), and may incorporate ATM cells or a variable-length data PDU. The variable-length PDU includes a pair of 48-bit addresses, data, and a CRC sum. In special cases, the MAC Header may encapsulate multiple MAC frames (see Section 6.2.5.4).

6.1.3 Future Use

A number of fields are defined as being “for future use” in the various MAC frames described in this document. These fields **MUST NOT** be interpreted or used in any manner by this version (1.0) of the MAC protocol.

6.2 MAC Frame Formats

6.2.1 Generic MAC Frame Format

A MAC frame is the basic unit of transfer between MAC sublayers at the CMTS and the cable modem. The same basic structure is used in both the upstream and downstream directions. MAC frames are variable in length. The term “frame” is used in this context to indicate a unit of information that is passed between MAC sublayer peers. This is not to be confused with the term “framing” that indicates some fixed timing relationship.

There are three distinct regions to consider, as shown in Figure 6-1. Preceding the MAC frame is either PMD sublayer overhead (upstream) or an MPEG transmission convergence header (downstream). The first part of the MAC frame is the MAC Header. The MAC Header uniquely identifies the contents of the MAC frame. Following the header is the optional Data PDU region. The format of the Data PDU and whether it is even present is described in the MAC Header.




Figure 6-1. Generic MAC Frame Format

6.2.1.1 PMD Overhead

In the upstream direction, the PHY layer indicates the start of the MAC frame to the MAC sublayer. From the MAC sublayer’s perspective, it only needs to know the total amount of overhead so it can account for it in the Bandwidth Allocation process. More information on this may be found in the PMD Sublayer section of this document (Section 6.4).

The FEC overhead is spread throughout the MAC frame and is assumed to be transparent to the MAC data stream. The MAC sublayer does need to be able to account for the overhead when doing Bandwidth Allocation. More information on this may be found in the Upstream Bandwidth Allocation section of this document (Section 6.4).

6.2.1.2 MAC Frame Transport

The transport of MAC frames by the PMD sublayer for upstream channels is shown in Figure 6-2.




Figure 6-2. Upstream MAC/PMD Convergence

The layering of MAC frames over MPEG in the downstream channel is described in Section 5.

6.2.1.3 Ordering of Bits and Octets

Within an octet, the least-significant bit is the first transmitted on the wire. This follows the convention used by Ethernet and [ISO 8802-3]. This is often called bit-little-endian order⁴.

Within the MAC layer, when numeric quantities are represented by more than one octet (i.e., 16-bit and 32-bit values), the octet containing the most-significant bits is the first transmitted on the wire. This section follows the textual convention that when bit-fields are presented in tables, the most-significant bits are topmost in the table. For example, in Table 6-2, FC_TYPE occupies the two most-significant bits and EHDR_ON occupies the least-significant bit. This is sometimes called byte-big-endian order.

6.2.1.3.1 Representing Negative Numbers

Signed integer values will be transmitted and received in two’s complement format.

6.2.1.3.2 Type-Length-Value Fields

Many MAC messages incorporate Type-Length-Value (TLV) fields. TLV fields MAY be unordered lists of TLV-tuples. Some TLV’s MAY be nested (see Appendix C). All TLV Length fields MUST be greater than zero.

6.2.1.4 MAC Header Format

The MAC Header format MUST be as shown in Figure 6-3.




Figure 6-3. MAC Header Format

All MAC Headers MUST have the general format as shown in Table 6-1. The Frame Control (FC) field is the first byte and uniquely identifies the rest of the contents within the MAC Header. The FC field is followed by 3 bytes of MAC control; an OPTIONAL Extended Header field (EHDR); plus a Header Check Sequence (HCS) to ensure the integrity of the MAC Header.

⁴ This applies to the upstream channel only. For the downstream channel, the MPEG transmission convergence sublayer presents an octet-wide interface to the MAC, so the MAC sublayer does not define the bit order.

Table 6-1. Generic MAC Header Format

MAC Header Field	Usage	Size
FC	Frame Control: Identifies type of MAC Header	8 bits
MAC_PARM	Parameter field whose use is dependent on FC: if EHDR_ON=1; used for EHDR field length (ELEN) else if for concatenated frames (see Table 6-9) used for MAC frame count else (for Requests only) indicates the number of mini-slots and/or ATM cells requested	8 bits
LEN (SID)	The length of the MAC frame. The length is defined to be the sum of the number of bytes in the extended header (if present) and the number of bytes following the HCS field. (For a REQ Header, this field is the Service ID instead)	16 bits
EHDR	Extended MAC Header (where present; variable size).	0-240 bytes
HCS	MAC Header Check Sequence	2 bytes
	Length of a MAC Header	6 bytes + EHDR

The HCS field is a 16-bit CRC that ensures the integrity of the MAC Header, even in a collision environment. The HCS field coverage MUST include the entire MAC Header, starting with the FC field and including any EHDR field that may be present. The HCS is calculated using CRC-CCITT ($x^{16} + x^{12} + x^5 + 1$) as defined in [ITU-T X.25].

The FC field is broken down into the FC_TYPE sub-field, FC_PARM sub-field and an EHDR_ON indication flag. The format of the FC field MUST be as shown in Table 6-2.

Table 6-2. FC Field Format

FC Field	Usage	Size
FC_TYPE	MAC Frame Control Type field: 00: Packet PDU MAC Header 01: ATM PDU MAC Header 10: Reserved PDU MAC Header 11: MAC Specific Header	2 bits
FC_PARM	Parameter bits, use dependent on FC_TYPE.	5 bits
EHDR_ON	When = 1, indicates that EHDR field is present. [Length of EHDR (ELEN) determined by MAC_PARM field]	1 bit

The FC_TYPE sub-field is the two MSBs of the FC field. These bits MUST always be interpreted in the same manner to indicate one of four possible MAC frame formats. These types include: MAC Header with Packet PDU; MAC Header with ATM cells; MAC Header reserved for future PDU types; or a MAC Header used for specific MAC control purposes. These types are spelled out in more detail in the remainder of this section.

The five bits following the FC_TYPE sub-field is the FC_PARM sub-field. The use of these bits are dependent on the type of MAC Header. The LSB of the FC field is the EHDR_ON indicator. If this bit is set, an Extended Header (EHDR) is present. The EHDR provides a mechanism to allow the MAC Header to be extensible in an inter-operable manner.

The Transmission Convergence Sublayer stuff-byte pattern is defined to be a value of 0xFF. This precludes the use of FC byte values which have FC_TYPE = ‘11’ and FC_PARM = ‘11111’.

The MAC_PARM field of the MAC Header serves several purposes depending on the FC field. If the EHDR_ON indicator is set, the MAC_PARM field MUST be used as the Extended Header length (ELEN). The EHDR field MAY vary from 0 to 240 bytes. If this is a concatenation MAC Header, the MAC_PARM field represents the number of MAC frames (CNT) in the concatenation (see Section 6.2.5.4). If this is a Request MAC Header (REQ) (see Section

6.2.5.3), the MAC_PARM field represents the amount of bandwidth being requested. In all other cases, the MAC_PARM field is reserved for future use.

The third field has two possible uses. In most cases, it indicates the length (LEN) of this MAC frame. In one special case, the Request MAC Header, it is used to indicate the cable modem's Service ID since no PDU follows the MAC Header.

The Extended Header (EHDR) field provides extensions to the MAC frame format. It is used to implement data link security and can be extended to add support for additional functions in future releases. Initial implementations SHOULD pass this field to the processor. This will allow future software upgrades to take advantage of this capability. (Refer to Section 6.2.6, "Extended MAC Headers" for details.)

6.2.1.5 Data PDU

The MAC Header MAY be followed by a Data PDU. The type and format of the Data PDU is defined in the Frame Control field of the MAC Header. The FC field explicitly defines a Packet Data PDU, an ATM Data PDU, a MAC Header only frame (no PDU) and a reserved code point (used as an escape mechanism for future extensions). All CMs MUST use the length in the MAC Header to skip over any reserved data.

6.2.2 Packet-Based MAC Frames

6.2.2.1 Variable-Length Packets

The MAC sublayer MUST support a variable-length Ethernet/[ISO8802-3]-type Packet Data PDU. The Packet PDU MUST be passed across the network in its entirety, including its original CRC. A unique Packet MAC Header is appended to the beginning. The frame format without an Extended header MUST be as shown in Figure 6-4 and Table 6-3.




Figure 6-4. Ethernet/802.3 Packet PDU Format

Table 6-3. Packet PDU Format

Field	Usage	Size
FC	FC_TYPE = 00; Packet MAC Header FC_PARM[4:0] = 00000; other values reserved for future use and ignored EHDR_ON = 0; no EHDR present this example	8 bits
MAC_PARM	Reserved, MUST be set to zero if there is no EHDR; otherwise set to length of EHDR	8 bits
LEN	LEN = n; length of Packet PDU in bytes	16 bits
EHDR	Extended MAC Header not present in this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	Packet PDU: DA - 48 bit Destination Address SA - 48 bit Source Address Type/Len - 16 bit Ethernet Type or [ISO8802-3] Length Field User Data (variable length, 0-1500 bytes) CRC - 32-bit CRC over packet PDU (as defined in Ethernet/[ISO8802-3])	n bytes
	Length of Packet MAC frame	6 + n bytes

6.2.3 ATM Cell MAC Frames

ATM transport is not defined in this specification.

In order to allow current frame-based CMs to operate in a possible future downstream channel in which ATM cells and frames are mixed, a codepoint for ATM has been defined. This will allow current modems to ignore ATM cells while receiving frames. The frame format MUST be as shown in Figure 6-5 and Table 6-4.




Figure 6-5. ATM Cell MAC Frame Format

Table 6-4. ATM Cell MAC Frame Format

Field	Usage	Size
FC	FC_TYPE = 01; ATM cell format MAC Header FC_PARM[4:0] = 00000; other values reserved for future use and ignored EHDR_ON = 0; no EHDR present this example	8 bits
MAC_PARM	Reserved, MUST be set to zero if there is no EHDR; otherwise set to length of EHDR	8 bits
LEN	LEN = nx53; length of ATM cell PDU, in bytes	16 bits
EHDR	Extended MAC Header not present this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
ATM Data	ATM cell PDU	n x 53 bytes
	Length of ATM cells based MAC frame	6 + n x 53 bytes

6.2.4 Reserved PDU MAC Frames

The MAC sublayer provides a reserved FC code point to allow for support of future (to be defined) PDU formats. The FC field of the MAC Header indicates that a Reserved PDU is present. This PDU MUST be silently discarded by MAC implementations of this version (1.0) of the specification. Compliant version 1.0 implementations MUST use the length field to skip over the Reserved PDU.

The format of the Reserved PDU without an extended header MUST be as shown in Figure 6-6 and Table 6-5.




Figure 6-6. Reserved PDU Format

Table 6-5. Reserved PDU Format

Field	Usage	Size
FC	FC_TYPE = 10; Reserved PDU MAC Header FC_PARM[4:0]; reserved for future use EHDR_ON = 0; no EHDR present this example	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of Reserved PDU in bytes	16 bits
EHDR	EHDR = 0; Extended MAC Header not present this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
User Data	Reserved Data PDU	n bytes
	Length of a Reserved PDU MAC frame	6 + n bytes

6.2.5 MAC-Specific Headers

There are several MAC Headers which are used for very specific functions. These functions include support for downstream timing and upstream ranging/power adjust, requesting bandwidth and concatenating multiple MAC frames.

6.2.5.1 Timing Header

A specific MAC Header is identified to help support the timing and adjustments required. In the downstream, this MAC Header MUST be used to transport the Global Timing Reference to which all cable modems synchronize. In the upstream, this MAC Header MUST be used as part of the Ranging message needed for a cable modem’s timing and power adjustments. The Timing MAC Header is followed by a Packet Data PDU. The format MUST be as shown in Figure 6-7 and Table 6-6.




Figure 6-7. Timing MAC Header

Table 6-6. Timing MAC Header Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 00000; Timing MAC Header EHDR_ON = 0; extended header prohibited for SYNC and RNG-REQ	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of Packet PDU in bytes	16 bits
EHDR	Extended MAC Header not present	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	MAC Management message: SYNC message (downstream only) RNG-REQ (upstream only)	n bytes
	Length of Timing Message MAC frame	6 + n bytes

6.2.5.2 MAC Management Header

A specific MAC Header is identified to help support the MAC management messages required. This MAC Header MUST be used to transport all MAC management messages (refer to Section 6.3). The format MUST be as shown Figure 6-8 and Table 6-7.




Figure 6-8. Management MAC Header

Table 6-7. Management MAC Header Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 00001 EHDR_ON	8 bits
MAC_PARM	Reserved for future use	8 bits
LEN	LEN = n; length of Packet PDU in bytes	16 bits
EHDR	Extended MAC Header not present this example	0 bytes
HCS	MAC Header Check Sequence	2 bytes
Packet Data	MAC Management message:	n bytes
	Length of Management MAC frame	6 + n bytes + EHDR

6.2.5.3 Request MAC Header

The Request MAC Header is the basic mechanism that a cable modem uses to request bandwidth. As such, it is only applicable in the upstream. There MUST be no Data PDUs following the Request MAC Header. The general format of the Request MUST be as shown in Figure 6-9 and Table 6-8.




Figure 6-9. Request MAC Header Format

Table 6-8. Request MAC Header (REQ) Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-Specific Header FC_PARM[4:0] = 00010; MAC Header only; no data PDU following EHDR_ON = 0; no EHDR allowed	8 bits
MAC_PARM	REQ, total amount of bandwidth requested (upstream only): if FC_PARM [4] = 0; REQ is number of mini-slots if FC_PARM [4] = 1; REQ is number of ATM cells	8 bits
SID	Service ID (0...0x1FFF)	16 bits
EHDR	Extended MAC Header not allowed	0 bytes
HCS	MAC Header Check Sequence	2 bytes
	Length of a REQ MAC Header	6 bytes

Because the Request MAC Header does not have a Data PDU following it, the LEN field is not needed. The LEN field MUST be replaced with an SID. The SID MUST uniquely identify a particular service queue within a given station.

The bandwidth request, REQ, MUST be specified in mini slots. The REQ field MUST indicate the current total amount of bandwidth requested for this service queue including appropriate allowance for the PHY overhead.

6.2.5.4 Concatenation

A Specific MAC Header is defined to allow multiple MAC frames to be concatenated. This allows a single MAC “burst” to be transferred across the network. The PHY overhead and the Concatenation MAC Header only occur once. Concatenation of multiple MAC frames MUST be as shown in Figure 6-10.

A compliant CMTS & CM MAY support concatenation.

Note: Concatenation only applies to upstream traffic. Concatenation MUST NOT be used on downstream traffic.




Figure 6-10. Concatenation of Multiple MAC Frames

Only one Concatenation MAC Header MUST be present per MAC “burst.” Nested concatenation MUST NOT be allowed. Immediately following the Concatenation MAC Header MUST be the MAC Header of the first MAC frame. Information within the MAC Header indicates the length of the first MAC Frame and provides a means to find the start of the next MAC Frame. Each MAC frame within a concatenation MUST be unique and MAY be of any type. This means that Packet, ATM, Reserved PDU and MAC-specific Frames MAY be mixed together. The embedded MAC frames MAY be addressed to different destinations and MUST be delivered as if they were transmitted individually.

The format of the Concatenation MAC Header MUST be as shown in Figure 6-11 and Table 6-9.




Figure 6-11. Concatenation MAC Header Format

Table 6-9. Concatenated MAC Frame Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 11100; Concatenation MAC Header EHDR_ON = 0; no EHDR with Concatenation Header	8 bits
MAC_PARM	CNT, number of MAC frames in this concatenation CNT = 0 indicates unspecified number of MAC frames	8 bits
LEN	LEN = x + ... + y; length of all following MAC frames in bytes	16 bits
EHDR	Extended MAC Header MUST NOT be used	0 bytes
HCS	MAC Header Check Sequence	2 bytes
MAC frame 1	first MAC frame: MAC Header plus OPTIONAL data PDU	x bytes
MAC frame n	last MAC frame: MAC Header plus OPTIONAL data PDU	y bytes
	Length of Concatenated MAC frame	6 + LEN bytes

The MAC_PARM field in the Concatenation MAC header provides a count of MAC frames as opposed to EHDR length or REQ amount as used in other MAC headers. If the field is non-zero, it MUST indicate the total count of MAC Frames (CNT) in this concatenation burst.

6.2.6 Extended MAC Headers

Every MAC Header, except the Timing, Concatenation MAC Header and Request Frame, has the capability of defining an Extended Header field (EHDR). The presence of an EHDR field MUST be indicated by the EHDR_ON flag in the FC field being set. Whenever this bit is set, the MAC_PARM field MUST be used as the EHDR length (ELEN). The minimum defined EHDR is 1 byte. The maximum EHDR length is 240 bytes.

A compliant CMTS & CM MUST support extended headers.

The format of a generic MAC Header with an Extended Header included MUST be as shown in Figure 6-12 and Table 6-10.

Note: Extended Headers MUST NOT be used in a Concatenation MAC Header, but MAY be included as part of the MAC Headers within the concatenation.

Extended Headers MUST NOT be used in Request and Timing MAC Headers.




Figure 6-12. Extended MAC Format

Table 6-10. Extended Header Format

Field	Usage	Size
FC	FC_TYPE = XX; Applies to all MAC Headers FC_PARM[4:0] = XXXXX; dependent on FC_TYPE EHDR_ON = 1; EHDR present this example	8 bits
MAC_PARM	ELEN = x; length of EHDR in bytes	8 bits
LEN	LEN = x + y; length of EHDR plus OPTIONAL data PDU in bytes	16 bits
EHDR	Extended MAC Header present this example	x bytes
HCS	MAC Header Check Sequence	2 bytes
PDU	OPTIONAL data PDU	y bytes
	Length of MAC frame with EHDR	6 + x + y bytes

Since the EHDR increases the length of the MAC frame, the LEN field MUST be increased to include both the length of the Data PDU and the length of the EHDR.

The EHDR field consists of one or more EH elements. Each EH element is variable sized. The first byte of the EH element MUST contain a type and a length field. Every CM MUST use this length to skip over any unknown EH elements. The format of an EH element MUST be as shown in Table 6-11.

Table 6-11. EH Element Format

EH Element Fields	Usage	Size
EH_TYPE	EH element Type Field	4 bits
EH_LEN	Length of EH_VALUE	4 bits
EH_VALUE	EH element data	0-15 bytes

The types of EH element defined in Table 6-12 MUST be supported. Reserved and extended types are undefined at this point and SHOULD be ignored.

The first eight EH element types are intended for one-way transfer between the cable modem and the CMTS. The next seven EH element types are for end-to-end usage within a MAC-sublayer domain. Thus, the information attached to the EHDR on the upstream MUST also be attached when the information is forwarded. The final EH element type is an escape mechanism that allows for more types and longer values, and MUST be as shown in Table 6-12.

Table 6-12. EH Element Format

EH_TYPE	EH_LEN	EH_VALUE
0	0	Null configuration setting; may be used to pad the extended header. The EH_LEN MUST be zero, but the configuration setting may be repeated.
1	3	Request: mini-slots requested (1 byte); SID (2 bytes) [CM --> CMTS]
2	2	Acknowledgment requested; SID (2 bytes) [CM --> CMTS]
3 (= BP_UP)	4	Upstream Privacy EH element [DOCSIS8]
4 (= BP_DOWN)	4	Downstream Privacy EH element [DOCSIS8]
5		Reserved
6		Reserved
7		Reserved
8		Reserved
10-14		Reserved [CM <-> CM]
15	XX	Extended EH element: EHX_TYPE (1 byte), EHX_LEN (1 byte), EH_VALUE (length determined by EHX_LEN)

6.2.7 Error-Handling

The cable network is a potentially harsh environment that may cause several different error conditions to occur. This section, together with Section 7.2.15, describes the procedures that are required when an exception occurs at the MAC framing level.

The most obvious type of error occurs when the HCS on the MAC Header fails. This may be a result of noise on the network or possibly by collisions in the upstream channel. Framing recovery on the downstream channel is performed by the MPEG transmission convergence sublayer. In the upstream channel, framing is recovered on each transmitted burst, such that framing on one burst is independent of framing on prior bursts. Hence, framing errors within a burst are handled by simply ignoring that burst; i.e., errors are unrecoverable until the next burst.

A second exception, which applies only to the upstream, occurs when the Length field is corrupted and the MAC thinks the frame is longer or shorter than it actually is. Synchronization will recover at the next valid upstream data interval.

For every MAC transmission, The HCS MUST be verified. When a bad HCS is detected, the MAC Header and any payload MUST be dropped. For Packet PDU transmissions, a bad CRC MAY be detected. Since the CRC only covers the Data PDU and the HCS covers the MAC Header; the MAC Header is still considered valid. Thus, the Packet PDU MUST be dropped, but any pertinent information in the MAC Header (e.g., bandwidth request information) MAY be used.

6.2.7.1 Error Codes and Messages

Appendix I lists CM and CMTS error codes and messages. These codes MUST be used as indicated in [DOCSIS5] and MAY be used for reporting errors via vendor-specific interfaces. If the error codes are used the error messages MAY be replaced by other descriptive messages.

6.3 MAC Management Messages

6.3.1 MAC Management Message Header

MAC management messages MUST be encapsulated in an LLC unnumbered information frame per [ISO8802-2], which in turn is encapsulated within the cable network MAC framing, as shown in Figure 6-13. Figure 6-13 shows the MAC Header and the MAC Management Message Header fields which are common across all MAC Management Messages.




Figure 6-13. MAC Header and MAC Management Message Header Fields

The fields MUST be as defined below.

- FC, MAC PARM, LEN, HCS** Common MAC frame header: refer to Section 6.2.1.4 for details. All messages use a MAC-specific header.
- Destination Address (DA)** MAC management frames will be addressed to a specific CM unicast address or to the DOCSIS management multicast address. These DOCSIS MAC management addresses are described in Appendix A.
- Source Address (SA)** The MAC address of the source CM or CMTS system.
- Msg Length** Length of the MAC message from DSAP to the end of the payload.
- DSAP** The LLC null SAP (00) as defined by [ISO8802-2].
- SSAP** The LLC null SAP (00) as defined by [ISO8802-2].
- Control** Unnumbered information frame (03) as defined by [ISO8802-2].
- Version** 1 octet
This field defines the version of the MAC management protocol in use. Set to 1 for this version.

Type 1 octet
This field defines the type of this particular MAC management message

Table 6-13. MAC Management Message Types

Type Value	Message Name	Message Description
1	SYNC	Timing Synchronization
2	UCD	Upstream Channel Descriptor
3	MAP	Upstream Bandwidth Allocation
4	RNG-REQ	Ranging Request
5	RNG-RSP	Ranging Response
6	REG-REQ	Registration Request
7	REG-RSP	Registration Response
8	UCC-REQ	Upstream Channel Change Request
9	UCC-RSP	Upstream Channel Change Response
10	TRI-TCD	Telephony Channel Descriptor [DOCSIS6]
11	TRI-TSI	Termination System Information [DOCSIS6]
12	BPKM-REQ	Privacy Key Management Request [DOCSIS8]
13	BPKM-RSP	Privacy Key Management Response [DOCSIS8]
14-255		Reserved for future use

RSVD 1 octet
This field is used to align the message payload on a 32 bit boundary.
Set to 0 for this version.

Management Message Payload variable length
As defined for each specific management message.

CRC Covers message including header fields (DA, SA,...). Polynomial defined by [ISO8802-3].

6.3.2 MAC Management Messages

A compliant CMTS or CM MUST support the following management message types.

6.3.2.1 Time Synchronization (SYNC)

Time Synchronization (SYNC) MUST be transmitted by CMTS at a periodic interval to establish MAC sublayer timing. This message MUST use an FC field of type: Timing. This MUST be followed by a Packet PDU in the format shown in Figure 6-14.




Figure 6-14. Format of Packet PDU Following the Timing Header

The parameters shall be as defined below.

CMTS Timestamp The count state of an incrementing 32 bit binary counter clocked with the CMTS 10.24 Mhz master clock. The CMTS timestamp represents the count state at the instant that the first byte (or a fixed time offset from the first byte) of the Time Synchronization MAC Management Message is transferred from the Downstream Transmission Convergence Sublayer to the Downstream Physical Media Dependent Sublayer as described in Section 4.3.7. The CMTS MUST NOT allow SYNC messages to cross MPEG packet boundary⁵).

6.3.2.2 Upstream Channel Descriptor (UCD)

An Upstream Channel Descriptor MUST be transmitted by the CMTS at a periodic interval to define the characteristics of an upstream channel (Figure 6-15). A separate message MUST be transmitted for each active upstream.

To provide for flexibility the message parameters following the channel ID MUST be encoded in a type/length/ value (TLV) form in which the type and length fields are each 1 octet long. Using this encoding, new parameters MAY be added which not all CMs can interpret. A CM which does not recognize a parameter type MUST skip over this parameter and MUST NOT treat the event as an error condition.




Figure 6-15. Upstream Channel Descriptor

⁵ Since the SYNC message applies to all upstream channels within this MAC domain, units were chosen to be independent of the symbol rate of any particular upstream channel. A timebase tick represents the smallest possible mini-slot at the highest possible symbol rate. See Section 6.5.4 for time-unit relationships.

A CMTS MUST generate UCDs in the format shown in Figure 6-15, including all of the following parameters:

Configuration Change Count Incremented by one (modulo the field size) by the CMTS whenever any of the values of this channel descriptor change. If the value of this count in a subsequent UCD remains the same, the CM can quickly decide that the remaining fields have not changed, and may be able to disregard the remainder of the message. This value is also referenced from the MAP.

Mini-Slot Size The size T of the Mini-Slot for this upstream channel in units of the Timebase Tick of 6.25 μ s. Allowable values are $T = 2^M$, $M = 1, \dots, 7$. That is, T = 2, 4, 8, 16, 32, 64 or 128.

upstream channel ID The identifier of the upstream channel to which this message refers. This identifier is arbitrarily chosen by the CMTS and is only unique within the MAC-Sublayer domain.

downstream channel ID The identifier of the downstream channel on which this message has been transmitted. This identifier is arbitrarily chosen by the CMTS and is only unique within the MAC-Sublayer domain.

All other parameters are coded as TLV tuples. The type values used MUST be those defined in 6-14 for channel parameters and Table 6-15 for upstream physical-layer burst attributes. Channel-wide parameters (types 1-3 in 6-14) MUST precede burst descriptors (type 4 below).

Table 6-14. Channel TLV Parameters

Name	Type (1 byte)	Length (1 byte)	Value (Variable length)
Symbol Rate	1	1	Multiples of base rate of 160 ksym/sec. (Value is 1, 2, 4, 8, or 16.)
Frequency	2	4	Upstream center frequency (Hz)
Preamble Pattern	3	1-128	Preamble superstring. All burst-specific preamble values are chosen as bit-substrings of this string. The first byte of the Value field contains the first 8 bits of the superstring, with the first bit of the preamble superstring in the MSB position of the first Value field byte, the eighth bit of the preamble superstring in the LSB position of the first Value field byte; the second byte in the Value field contains the second eight bits of the superstring, with the ninth bit of the superstring in the MSB of the second byte and sixteenth bit of the preamble superstring in the LSB of the second byte, and so forth.
Burst Descriptor	4	n	May appear more than once; described below.

Burst Descriptors are composed of an upstream Interval Usage Code, followed by a TLV encodings that define, for each type of upstream usage interval, the physical-layer characteristics that are to be used during that interval. The upstream interval usage codes are defined in the MAP message (see Section 6.3.2.3 and Table 6-16). The format of the Burst Descriptor is shown in Figure 6-16.




Figure 6-16. Top-Level Encoding for a Burst Descriptor

In Figure 6-16:

Type	4 for Burst Descriptor.
Length	The number of bytes in the overall object, including the IUC and the embedded TLV items.
IUC	Interval Usage code defined in Table 6-16. The IUC is coded on the 4 less significant bits. The 4 most significant bits are unused (=0).
TLV items	TLV parameters described in Table 6-15.

A Burst Descriptor **MUST** be included for each interval usage code that is to be used in the allocation MAP. The interval usage code above must be one of the values from Table 6-16.

Within each Burst Descriptor is an unordered list of Physical-layer attributes, encoded as TLV values. These attributes are shown in Table 6-15.

Table 6-15. Upstream Physical-Layer Burst Attributes

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Modulation Type	1	1	1 = QPSK, 2 = 16QAM
Differential Encoding	2	1	1 = on, 2 = off
Preamble Length	3	2	Up to 1024 bits. The value must be an integral number of symbols (a multiple of 2 for QPSK and 4 for 16QAM)
Preamble Value Offset	4	2	Identifies the bits to be used for the preamble value. This is specified as a starting offset into the Preamble Pattern (see Table 6-14). That is, a value of zero means that the first bit of the preamble for this burst type is the value of the first bit of the Preamble Pattern. A value of 100 means that the preamble is to use the 101st and succeeding bits from the Preamble Pattern. This value must be a multiple of the symbol size. The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is I1 in the first symbol of the burst (see Section 4.2.2.2).
FEC Error Correction (T bytes)	5	1	0-10 bytes. Zero implies no Forward Error Correction.
FEC Codeword Information Bytes (k)	6	1	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on) (Not used if no FEC, T=0)
Scrambler Seed	7	2	The 15-bit seed value left-justified in the 2-byte field. Bit 15 is the MSB of the first byte and the LSB of the second byte is not used. (Not used if scrambler is off)
Maximum Burst Size	8	1	The maximum number of mini-slots that can be transmitted during this burst type. Absence of this configuration setting implies that the burst size is limited elsewhere. When the interval type is Short Data Grant this value MUST be present and greater than zero. (See 6.4.1.1.5)
Guard Time Size	9	1	Number of symbol times which must follow the end of this burst. (Although this value may be derivable from other network and architectural parameters, it is included here to ensure that the CMs and CMTS all use the same value.)
Last Codeword Length	10	1	1 = fixed; 2 = shortened
Scrambler on/off	11	1	1 = on; 2 = off

6.3.2.2.1 Example of UCD Encoded TLV Data

An example of UCD encoded TLV data is given in Figure 6-17.




Figure 6-17. Example of UCD Encoded TLV Data

6.3.2.3 Upstream Bandwidth Allocation Map (MAP)

A CMTS MUST generate MAPs in the format shown in Figure 6-18.




Figure 6-18. MAP Format

The parameters MUST be as follows:

- Upstream Channel ID** The identifier of the upstream channel to which this message refers.
- UCD Count** Matches the value of the Configuration Change Count of the UCD which describes the burst parameters which apply to this map. See Section 7.2.13.
- Number Elements** Number of information elements in the map.
- Reserved** Reserved field for alignment.
- Alloc Start Time** Effective start time from CMTS initialization (in mini-slots) for assignments within this map.
- Ack Time** Latest time, from CMTS initialization, (mini-slots) processed in upstream. This time is used by the CMs for collision detection purposes. See Section 6.4.4.
- Ranging Backoff Start** Initial back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest-order bits must be unused and set to 0).
- Ranging Backoff End** Final back-off window for initial ranging contention, expressed as a power of two. Values range 0-15 (the highest-order bits must be unused and set to 0).
- Data Backoff Start** Initial back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest-order bits must be unused and set to 0).
- Data Backoff End** Final back-off window for contention data and requests, expressed as a power of two. Values range 0-15 (the highest-order bits must be unused and set to 0).
- MAP information elements** MUST be in the format defined in Figure 6-19 and Table 6-16. Values for UCDs are defined in Table 6-16 and are described in detail in Section 6.4.1.

Note: The lower (26-M) bits of the Alloc Start Time and Ack Time MUST be used as the effective MAP start and act times where M is given in Section 6.3.2.2. The relationship between the Alloc Start/Ack time counters and the timestamp counter is described in Section 6.5.4.




Figure 6-19. MAP Information Element Structure

Table 6-16. Allocation MAP Information Elements (IE)

IE Name	Interval Usage Code (IUC) (4 bits)	SID (14 bits)	Mini-slot Offset (14 bits)
Request	1	any	starting offset of REQ region
REQ/Data (refer to Appendix A for multicast definition)	2	multicast	starting offset of IMMEDIATE Data region well-known multicasts define start intervals
Initial Maintenance	3	broadcast/ multicast	starting offset of MAINT region (used in Initial Ranging)
Station Maintenance ^a	4	unicast	starting offset of MAINT region (used in Periodic Ranging)
Short Data Grant ^b	5	unicast	starting offset of Data Grant assignment; if inferred length = 0, it is a Data Grant pending.
Long Data Grant	6	unicast	starting offset of Data Grant assignment; if inferred length = 0, it is a Data Grant Pending
Null IE	7	zero	ending offset of the previous grant. Used to bound the length of the last actual interval allocation.
Data Ack	8	unicast	CMTS sets to map length
Reserved	9-14	any	reserved
Expansion	15	expanded IUC	# of additional 32-bit words in this IE

- a. Although the distinction between Initial Maintenance and Station Maintenance is unambiguous from the Service ID type, separate codes are used to ease physical-layer configuration (see burst descriptor encodings, Table 6-15).
 - b. The distinction between long and short data grants is related to the amount of data that can be transmitted in the grant. A short data grant interval may use FEC parameters that are appropriate to short packets while a long data grant may be able to take advantage of greater FEC coding efficiency.
- Each IE is a 32-bit quantity, of which the most significant 14 bits represent the SID, the middle 4 bits the IUC, and the low-order 14 bits the mini-slot offset.
- The SID used in the Station Maintenance IE **MUST** be a Temporary SID, or the first Registration SID (and maybe the only one) that was assigned in the REG-RSP message to a CM.

6.3.2.4 Ranging Request (RNG-REQ)

A Ranging Request **MUST** be transmitted by a CM at initialization and periodically on request from CMTS to determine network delay. This message **MUST** use an FC field of type: Timing. This **MUST** be followed by a Packet PDU in the format shown in Figure 6-20.




Figure 6-20. Packet PDU Following the Timing Header

Parameters MUST be as follows:

SID

For RNG-REQ messages transmitted in Initial Maintenance intervals:

- Initialization SID if modem is attempting to join the network.
- Initialization SID if modem has not yet registered and is changing upstream, downstream or both upstream and downstream channels as directed by a downloaded parameter file.
- Primary SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels.
- Registration SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels.

For RNG-REQ messages transmitted in Station Maintenance Intervals:

- Temporary SID if during or before registration.
- Primary SID if after registration.

This is a 16-bit field of which the lower 14 bits define the SID with bits 14,15 defined to be 0.

Downstream channel ID

The identifier of the downstream channel on which the CM received the UCD which described this upstream. This is an 8-bit field.

Pending Till Complete

If zero, all previous Ranging Response attributes have been applied prior to transmitting this request. If nonzero, this is the time estimated to be needed to complete assimilation of ranging parameters. Note that only equalization can be deferred. Units are in unsigned centiseconds (10 msec).

6.3.2.5 Ranging Response (RNG-RSP)

A Ranging Response MUST be transmitted by a CMTS in response to received RNG-REQ. The state machines describing the ranging procedure appear in Section 7.2.4. In that procedure it may be noted that, from the point of view of the CM, reception of a Ranging Response is stateless. In particular, the CM MUST be prepared to receive a Ranging Response at any time, not just following a Ranging Request.

To provide for flexibility, the message parameters following the Upstream Channel ID MUST be encoded in a type/length/value (TLV) form. Using this encoding, new parameters can be added which not all CMs can interpret. A CM which does not recognize a parameter type MUST simply skip over this parameter and MUST NOT treat the event as an error condition.




Figure 6-21. Ranging Response

A CMTS MUST generate Ranging Responses in the form shown in 6-21, including all of the following parameters:

SID If the modem is being instructed by this response to move to a different channel, this is the initialization SID. Otherwise, this is the SID from the corresponding RNG-REQ to which this response refers, except if the corresponding RNG-REQ was an initial ranging request specifying a initialization SID, in which case this is the assigned temporary SID.

Upstream channel ID The identifier of the upstream channel on which the CMTS received the RNG-REQ to which this response refers. On the first ranging response received by the CM during initial ranging, this channel ID may be different from the channel ID the CM used to transmit the range request (see Appendix G). Thus, the CM MUST use this channel ID for the rest of its transactions, not the channel ID it initiated the range request from.

Ranging status Used to indicate whether upstream messages are received within acceptable limits by CMTS.

All other parameters are coded as TLV tuples.

Timing adjust information The time by which to offset frame transmission so that frames arrive at the expected mini-slot time at the CMTS.

Power adjust information Specifies the relative change in transmission power level that the CM is to make in order that transmissions arrive at the CMTS at the desired power.

Frequency adjust information Specifies the relative change in transmission frequency that the CM is to make in order to better match the CMTS. (This is fine-frequency adjustment within a channel, not re-assignment to a different channel)

CM transmitter equalization information If the CM implements transmission equalization, this provides the equalization coefficients (optional).

Downstream Frequency Override An optional parameter. The downstream frequency with which the modem should redo initial ranging. (See Section 6.3.2.5.3)

Upstream Channel ID Override An optional parameter. The identifier of the upstream channel with which the modem should redo initial ranging. (See Section 6.3.2.5.3)

6.3.2.5.1 Encodings

The type values used MUST be those defined in Table 6-21 and Figure 6-22. These are unique within the ranging response message but not across the entire MAC message set. The type and length fields MUST each be 1 octet in length.

If the CM supports pre-equalization, the requirements in the remainder of this section related to pre-equalization MUST be supported. If the CM does not implement the pre-equalizer, the CM MUST ignore the Transmit Equalization Adjust TLV.

Table 6-21. Ranging Response Message Encodings

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Timing Adjust	1	4	TX timing offset adjustment (signed 32-bit, units of (6.25 microsec/64))
Power Level Adjust	2	1	TX Power offset adjustment (signed 8-bit, 1/4-dB units)
Offset Frequency Adjust	3	2	TX frequency offset adjustment (signed 16-bit, Hz units)
Transmit Equalization Adjust	4	n	TX equalization data - see details below
Ranging Status	5	1	1 = continue, 2 = abort, 3 = success
Downstream frequency override	6	4	Center frequency of new downstream channel in Hz
Upstream channel ID override	7	1	Identifier of the new upstream channel.
Reserved	8-255	n	Reserved for future use




Figure 6-22. Generalized Decision Feedback Equalization Coefficients

The number of forward taps per symbol MUST be either 1, 2, or 4. The main tap location refers to the position of the zero delay tap, between 1 and N. For a symbol-spaced equalizer, the number of forward taps per symbol field MUST

be set to “1”. The number of reverse taps (M) field MUST be set to “0” for a linear equalizer. The total number of taps MAY range up to 64. Each tap consists of a real and imaginary coefficient entry in the table.

If more than 255 bytes are needed to represent equalization information, then several type-4 elements MAY be used. Data MUST be treated as if byte-concatenated, that is, the first byte after the length field of the second type-4 element is treated as if it immediately followed the last byte of the first type-4 element.




Figure 6-23. Generalized Equalizer Tap Location Definition

6.3.2.5.2 Example of TLV Data

An example of TLV data is given in Figure 6-24.

Type 1	Length 4	Timing adjust
Type 2	Length 1	Power adjust
Type 3	Length 2	Frequency adjust information
Type 4	Length x	x bytes of CM transmitter equalization information
Type 5	Length 1	Ranging status

Figure 6-24. Example of TLV Data

6.3.2.5.3 Overriding Channels Prior to Registration

The RNG-RSP message allows the CMTS to instruct the modem to move to a new downstream and/or upstream channel and to repeat initial ranging. However, the CMTS may do this only in response to an initial ranging request from a modem that is attempting to join the network, or in response to any of the unicast ranging requests that take place immediately after this initial ranging and up to the point where the modem successfully completes periodic ranging. From this point on, only the C.7.4 and UCC-REQ mechanisms are available for moving the modem to a new upstream channel, and only the C.7.3 mechanism is available for moving the modem to a new downstream channel.

If a downstream frequency override is specified in the RNG-RSP, the modem MUST reinitialize its MAC (Section 7.2) using initial ranging with the specified downstream center frequency as the first scanned channel. For the upstream channel, the modem may select any valid channel based on received UCD messages.

If an upstream channel ID override is specified in the RNG-RSP, the modem MUST reinitialize its MAC (Section 7.2) using initial ranging with the upstream channel specified in the RNG-RSP for its first attempt and the same downstream frequency on which the RNG-RSP was received.

If both downstream frequency and upstream channel ID overrides are present in the RNG-RSP, the modem MUST reinitialize its MAC (Section 7.2) using initial ranging with the specified downstream frequency and upstream channel ID for its first attempt.

Note that when a modem with an assigned temporary SID is instructed to move to a new downstream and/or upstream channel and to redo initial ranging, the modem MUST consider the temporary SID to be deassigned. The modem MUST redo initial ranging using the Initialization SID.

Configuration file settings (Appendix C.3) for upstream channel ID and downstream frequency are optional, but if specified in the config file they take precedence over the ranging response parameters.

6.3.2.6 Registration Request (REG-REQ)

A Registration Request MUST be transmitted by a CM at initialization after receipt of a CM parameter file.

To provide for flexibility, the message parameters following the SID MUST be encoded in a type/length/value form. Using this encoding, new parameters MAY be added which not all CMTSes can interpret. A CMTS which does not recognize a parameter type MUST simply skip over this parameter and MUST not treat the event as an error condition.




Figure 6-25. Registration Request

A CM MUST generate Registration Requests in the form shown in Figure 6-25, including the following parameters:

SID Temporary SID for this CM

All other parameters are coded as TLV tuples as defined in Appendix C. Registration Requests can contain many different TLV parameters, some of which are set by the CM according to its configuration file and some of which are generated by the CM itself.

If found in the Configuration File, the following Configuration Settings MUST be included in the Registration Request.

Configuration File Settings:

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Network Access Configuration Setting
- Class of Service Configuration Setting
- Baseline Privacy Configuration Setting
- Vendor-Specific Information Configuration Setting
- Maximum Number of CPEs
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address
- CM MIC Configuration Setting
- CMTS MIC Configuration Setting

The following registration parameter **MUST** be included in the Registration Request.

Vendor Specific Parameter:

- Vendor ID Configuration Setting (Vendor ID of CM)

The following registration parameters **MAY** also be included in the Registration Request.

Other Registration Parameters:

- Modem Capabilities Configuration Setting
- Modem IP Address

The following Configuration Settings **MUST NOT** be forwarded to the CMTS in the Registration Request.

- Software Upgrade Filename
- Software Upgrade TFTP Server IP Address
- SNMP Write-Access Control
- SNMP MIB Object
- CPE Ethernet MAC Address
- End Configuration Setting
- Pad Configuration Setting
- Telephone Settings Option

6.3.2.6.1 Encodings

The type values used are unique within the registration request message but not across the entire MAC message set. They **MUST** be those defined in Appendix C.

Note: The CM **MUST** forward the vendor specific configuration settings to the CMTS in the same order in which they were received in the configuration file to allow the message integrity check to be performed.
Use of the end of data marker (type 255) and of the PAD (type 0) is not permitted in a registration request.

6.3.2.6.2 Example

An example of type value encodings is given in Figure 6-26.

Type 1	Length 4	Downstream Frequency	
Type 2	Length 1	Upstream channel	
Type 3	Length 1	Network Access	
Type 4	Length 28	Service class definition class 1	
Type 4	Length 28	Service class definition class 2	
Type 4	Length 28	Service class definition class n	
Type 5	Length 3	Modem capabilities	
Type 12	Length 4	Modem IP address	
Type 8	Length 3	Vendor ID	
Type 43	Length n	n bytes of vendor specific data	
Type 6	Length 16	CM message integrity check	
Type 6	Length 16	CMTS message integrity check	

Figure 6-26. Example of Registration Request Type Value Encodings

6.3.2.7 Registration Response (REG-RSP)

A Registration Response MUST be transmitted by CMTS in response to received REG-REQ.

To provide for flexibility, the message parameters following the Response field MUST be encoded in a TLV format. Using this encoding, new parameters MAY be added which not all CMs can interpret. A CM which does not recognize a parameter type MUST skip over this parameter and MUST NOT treat the event as an error condition.




Figure 6-27. Registration Response Format

A CMTS MUST generate Registration Responses in the form shown in Figure 6-27, including both of the following parameters:

SID from Corresponding

REG-REQ SID from corresponding REG-REQ to which this response refers

Response 0 = ok
 1 = authentication failure
 2 = class of service failure

All other parameters are coded TLV tuples

Modem Capabilities The CMTS response to the capabilities of the modem

Service Class Data Returned when Response = ok
 Service ID / service class tuple for each class of service granted

Service Not Available Returned when Response = class of service failure.
 If a service class cannot be supported, this configuration setting is returned in place of the service class data. If this is received, the entire registration request is considered to have failed and must be repeated.

Vendor-Specific Data As defined in Appendix C

- Vendor ID Configuration Setting (vendor ID of CMTS)
- Vendor-specific extensions

Note: Service class IDs MUST be those requested in the corresponding REG-REQ. The initialization SID MUST no longer be used once the REG-RSP is received.

6.3.2.7.1 Encodings

The type values used MUST be those shown below. These are unique within the registration response message but not across the entire MAC message set. The type and length fields MUST each be 1 octet.

6.3.2.7.1.1 Modem Capabilities

This field defines the CMTS response to the modem capability field in the Registration Request. The CMTS responds to the modem capabilities to indicate whether they may be used. If the CMTS does not recognize a modem capability, it must return this as “off” in the Registration Response.

Only capabilities set to “on” in the REG-REQ may be set “on” in the REG-RSP as this is the handshake indicating that they have been successfully negotiated.

Encodings are as defined for the Registration Request.

6.3.2.7.1.2 Service Class Data

This encoding defines the parameters associated with a requested class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class of service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated service class data configuration setting string. A single service class data configuration setting **MUST** be used to define the parameters for a single service class. Multiple class definitions **MUST** use multiple service class data configuration setting sets.

type	length	value
1	n	encoded service class data

Internal service class data encodings

Class ID

The value of the field **MUST** specify the identifier for the class of service to which the encapsulated string applies. This **MUST** be a class which was requested in the associated REG-REQ, if present.

Note: If no Class ID was present for any Class-of-Service TLV in the REG-REQ, the CMTS **MUST** send a REG-RSP with a class-of-service failure response and no service class TLVs.

Type	length	value
1	1	from REG-REQ

Valid Range

The class ID **MUST** be in the range 1 to 16.

Service ID

The value of the field **MUST** specify the SID associated with this service class.

Type	Length	Value
2	2	SID

6.3.2.7.2 Registration Response Encoding Example

An example of Registration Response encoding is given in Figure 6-28.

Type 1	Length 7	Service class definition class 1
Type 1	Length 7	Service class definition class 2
...		
Type 1	Length 7	Service class definition class n
Type 6	Length 6	Modem capability

Figure 6-28. Example of Registration Response Encoding

6.3.2.7.3 Sample Service Class Data Encoding

Sample service class data encodings are provided in Table 6-18.

Table 6-17. Sample Service Class Data Encoding

Type	Length	VALUE (sub)type	Length	Value	Description
1	7	1	1	1	service class data configuration setting service class 1
		2	2	123	SID for this class
1	7	1	1	2	service class data configuration setting service class 2
		2	2	244	SID for this class
1	7	1	1	n	service class data configuration setting service class n
		2	2	345	SID for this class

6.3.2.8 Upstream Channel Change Request (UCC-REQ)

An Upstream Channel Change Request MAY be transmitted by a CMTS to cause a CM to change the upstream channel on which it is transmitting. The format of an UCC-REQ message is shown in Figure 6-29.




Figure 6-29. Upstream Channel Change Request

Parameters MUST be as follows:

Upstream channel ID The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This is an 8-bit field.

6.3.2.9 Upstream Channel Change Response (UCC-RSP)

An Upstream Channel Change Response MUST be transmitted by a CM in response to a received Upstream Channel Change Request message to indicate that it has received and is complying with the UCC-REQ. The format of an UCC-RSP message is shown in Figure 6-30.

Before it begins to switch to a new upstream channel, a CM MUST transmit a UCC-RSP on its existing upstream channel. A CM MAY ignore an UCC-REQ message while it is in the process of performing a channel change. When a CM receives a UCC-REQ message requesting that it switch to an upstream channel that it is already using, the CM MUST respond with a UCC-RSP message on that channel indicating that it is already using the correct channel.

To switch to a new upstream channel, a CM will begin a new ranging procedure for that channel, and upon completion of ranging will proceed without re-performing registration. The full procedure for changing channels is described in Section 7.2.14.




Figure 6-30. Upstream Channel Change Response

Parameters MUST be as follows:

Upstream channel ID The identifier of the upstream channel to which the CM is to switch for upstream transmissions. This is the same Channel ID specified in the UCC-REQ message. This is an 8-bit field.

6.4 Upstream Bandwidth Allocation

The upstream channel is modeled as a stream of mini-slots. The CMTS MUST generate the time reference for identifying these slots. It MUST also control access to these slots by the cable modems. For example, it MAY grant some number of contiguous slots to a CM for it to transmit a data PDU. The CM MUST time its transmission so that the CMTS receives it in the time reference specified. This section describes the elements of protocol used in requesting, granting, and using upstream bandwidth. The basic mechanism for assigning bandwidth management is the allocation map. Please refer to Figure 6-31.

The allocation map is a MAC Management message transmitted by the CMTS on the downstream channel which describes, for some interval, the uses to which the upstream mini-slots MUST be put. A given map MAY describe some slots as grants for particular stations to transmit data in, other slots as available for contention transmission, and other slots as an opportunity for new stations to join the link.

Many different scheduling algorithms MAY be implemented in the CMTS by different vendors; this specification does not mandate a particular algorithm. Instead, it describes the protocol elements by which bandwidth is requested and granted.




Figure 6-31. Allocation Map

The bandwidth allocation MUST include the following basic elements:

- Each CM has one or more short (14-bit) service identifiers as well as a 48-bit address.
- Upstream bandwidth is divided into a stream of mini-slots. Each mini-slot is numbered relative to a master reference maintained by the CMTS. The clocking information is distributed to the CMs by means of SYNC packets.
- CMs MAY issue requests to the CMTS for upstream bandwidth.

The CMTS MUST transmit allocation map PDUs on the downstream channel defining the allowed usage of each mini-slot. The map is described below.

6.4.1 The Allocation Map MAC Management Message

The allocation map is a varying-length MAC Management message that is transmitted by the CMTS to define transmission opportunities on the upstream channel. It includes a fixed-length header followed by a variable number of information elements (IEs) in the format shown in Section 6.3.2.3. Each information element defines the allowed usage for a range of mini-slots.

Note that it should be understood by both CM and CMTS that the lower (26-M) bits of alloc start and ack times MUST be used as the effective MAP start and ack times, where M is defined in Section 6.3.2.2. The relationship between alloc start/ack time counters and the timestamp counter is further described in Section 6.5.4.

6.4.1.1 Information Elements

Each IE consists of a 14-bit Service ID, a 4-bit type code, and a 14-bit starting offset as defined in Section 6.3.2.3. Since all stations MUST scan all IEs, it is critical that IEs be short and relatively fixed format. IEs within the map are strictly ordered by starting offset. For most purposes, the duration described by the IE is inferred by the difference between the IE's starting offset and that of the following IE. For this reason, a null IE MUST terminate the list. Refer to Table 6-16.

Four types of Service IDs are defined:

1. 0x3FFF - broadcast, intended for all stations
2. 0x2000-0x3FFE - multicast, purpose is defined administratively. See Appendix A.
3. 0x0001-0x1FFF - unicast, intended for a particular CM or a particular service within that CM
4. 0x0000 - null address, addressed to no station.

6.4.1.1.1 The Request IE

The Request IE provides an upstream interval in which requests MAY be made for bandwidth for upstream data transmission. The character of this IE changes depending on the class of Service ID. If broadcast, this is an invitation for CMs to contend for requests. Section 6.4.4 describes which contention transmit opportunity may be used. If unicast, this is an invitation for a particular CM to request bandwidth. Unicasts MAY be used as part of a class-of-service implementation (see below). Packets transmitted in this interval MUST use the Request MAC Header format (see Section 6.2.5.3). This IE MUST be supported by a compliant CM or CMTS.

6.4.1.1.2 The Request/Data IE

The Request/Data IE provides an upstream interval in which requests for bandwidth or short data packets MAY be transmitted. This IE is distinguished from the Request IE in that:

- It provides a means by which allocation algorithms MAY provide for “immediate” data contention under light loads, and a means by which this opportunity can be withdrawn as network loading increases.
- Multicast Service IDs MUST be used to specify maximum data length, as well as allowed random starting points within the interval. For example, a particular multicast ID MAY specify a maximum of 64-byte data packets, with random starting points of every fourth slot.

A small number of well-known multicast Service IDs are defined in Appendix A. Others are available for vendor-specific algorithms.

Since data packets transmitted within this interval may collide, the CMTS MUST acknowledge any that are successfully received. The data packet MUST indicate in the MAC Header that a data acknowledgment is desired (see Table 6-12).

This IE MAY be supported by a compliant CM or CMTS.

6.4.1.1.3 The Initial Maintenance IE

The Initial Maintenance IE provides an interval in which new stations may join the network. A long interval, equivalent to the maximum round-trip propagation delay plus the transmission time of the Ranging Request (RNG-REQ) message (see Section 6.3.2.4), MUST be provided to allow new stations to perform initial ranging. Packets transmitted in this interval MUST use the RNG-REQ MAC Management message format (see Section 6.3.2.4).

6.4.1.1.4 The Station Maintenance IE

The Station Maintenance IE provides an interval in which stations are expected to perform some aspect of routine network maintenance, such as ranging or power adjustment. The CMTS MAY request that a particular CM perform some task related to network maintenance, such as periodic transmit power adjustment. In this case, the Station Maintenance IE is unicast to provide upstream bandwidth in which to perform this task. Packets transmitted in this interval MUST use the RNG-REQ MAC Management message format (see Section 6.3.2.4).

6.4.1.1.5 Short and Long Data Grant IEs

The Data Grant IE provides an opportunity for a CM to transmit one or more upstream PDUs. These IEs MAY also be used with an inferred length of zero mini slots (a zero length grant), to indicate that a request has been received and is pending (a Data Grant Pending). This IE is issued either in response to a request from a station, or because of an administrative policy providing some amount of bandwidth to a particular station (see class-of-service discussion below).

Short Data Grants are used with intervals less than or equal to the maximum burst size for this usage specified in the Upstream Channel Descriptor. If Short Data bursts are defined in the UCD, all Long Data Grants MUST be for a larger number of mini-slots than the maximum for Short Data. The distinction between Long and Short Data Grants may be exploited in physical-layer forward-error-correction coding; otherwise, it is not meaningful to the bandwidth allocation process.

If this IE is a Data Grant Pending (a zero length grant), it MUST follow the NULL IE. This allows cable modems to process all actual allocations first, before scanning the Map for request acknowledgments and data acknowledgments.

6.4.1.1.6 Data Acknowledge IE

The Data Acknowledge IE acknowledges that a data PDU was received. The CM MUST have requested this acknowledgment within the data PDU (normally this would be done for PDUs transmitted within a contention interval in order to detect collisions).

This IE MUST follow the NULL IE. This allows cable modems to process all actual interval allocations first, before scanning the Map for request acknowledgments and data acknowledgments.

6.4.1.1.7 Expansion IE

The Expansion IE provides for extensibility, if more than 16 code points or 32 bits are needed for future IEs.

6.4.1.1.8 Null IE

A Null IE terminates all actual allocations in the IE list. It is used to infer a length for the last interval. All Data Acknowledge IEs and All Data Grant Pending IEs (Data Grants with an inferred length of 0) must follow the Null IE

6.4.1.2 Requests and Grants

Only one type of upstream request is inherent to the allocation protocol: a request for upstream bandwidth. This request MAY be transmitted any time that either a request or a data PDU is allowed from the particular station. It MAY be transmitted during an interval described by any of:

- A Request IE
- A Request/Data IE⁶.
- A Data Grant IE.

In addition, it MAY be piggybacked⁷ on a data transmission. The request includes:

- The Service ID making the request
- The number of mini-slots or ATM cells requested

The number of mini-slots requested MUST be the total number that are desired by the CM at the time of the request (including any physical-layer overhead)⁸, subject to administrative limits⁹. The CM MUST request a number of mini-slots corresponding to one or more complete packets. A non-concatenating CM MUST request only the necessary mini-slots for one MAC frame per request. If, for whatever reason, a previous request has not been satisfied when the CM is making a new request, it MUST include the number of slots from the old request in the new total. Note that only one request at a time (per Service ID) will be outstanding. Because the CMTS MUST continue to issue null

⁶ The Request/Data IE MUST be used in conjunction with the Request IE since Request/Data IE support is a CM option.

⁷ When piggybacked, these values are carried in the Extended Header (Section 6.2.6, EH_TYPE=1).

⁸ Physical-layer overhead that MUST be accounted for in a request includes: guard band, preamble, and FEC which are dependent on the burst profile.

⁹ The CM is limited by the Maximum Transmit Burst for the service class, as defined in Appendix C.

grants for as long as a request is unsatisfied, the CM is able to unambiguously determine when its request is still pending.

Administrative limits MAY be assigned, either globally or per Service ID, on the number of mini-slots that MAY be requested at once. The global limit is configured as the maximum transmission burst size.

The CMTS MUST NOT make any data grant for more than 255 mini-slots (the maximum the CM is able to request). This limitation is necessary to allow CM implementations to avoid the costs associated with being prepared to handle arbitrarily long grants.

6.4.2 Map Transmission and Timing

The allocation map MUST be transmitted in time to propagate across the physical cable and be received and handled by the receiving CMs. As such, it MAY be transmitted considerably earlier than its effective time. The components of the delay are:

- Worst-case round-trip propagation delay—may be network-specific, but on the order of hundreds of microseconds.
- Queuing delays within the CMTS—implementation-specific.
- Processing delays within the CMs—MUST allow a minimum processing time by each CM as specified in B (CM MAP Processing Time).
- PMD-layer FEC interleaving.

Within these constraints, vendors MAY wish to minimize this delay so as to minimize latency of access to the upstream channel.

The number of mini-slots described MAY vary from map to map. At minimum, a map MAY describe a single mini-slot. This would be wasteful in both downstream bandwidth and in processing time within the CMs. At maximum, a map MAY stretch to tens of milliseconds. Such a map would provide poor upstream latency. Allocation algorithms MAY vary the size of the maps over time to provide a balance of network utilization and latency under varying traffic loads.

At minimum, a map MUST contain two Information Elements: one to describe an interval and a null IE to terminate the list. At a maximum, a map MUST be bounded by a limit of 240 information elements. Maps are also bounded in that they MUST NOT describe more than 4096 mini-slots into the future. The latter limit is intended to bound the number of future mini-slots that each CM is required to track. Even though several maps may be outstanding, the sum of the number of mini-slots they describe MUST NOT exceed 4096.

The set of all maps, taken together, MUST describe every mini-slot in the upstream channel. If a CM fails to receive a map describing a particular interval, it MUST NOT transmit during that interval.

Multiple maps MAY be outstanding at once.

6.4.3 Protocol Example

This section illustrates the interchange between the CM and the CMTS when the CM has data to transmit (Figure 6-32). Suppose a given CM has a data PDU available for transmission.




Figure 6-32. Protocol Example

Description

At time t_1 , the CMTS transmits a map whose effective starting time is t_3 . Within this map is a Request IE which will start at t_5 . The difference between t_1 and t_3 is needed to allow for:

- downstream propagation delay (including FEC interleaving) to allow all CMs to receive the Map
- processing time at the CM (allows the CMs to parse the Map and translate it into transmission opportunities)
- upstream propagation delay (to allow the CM's transmission of the first upstream data to begin in time to arrive at the CMTS at time t_3).

At t_2 , the CM receives this map and scans it for request opportunities. In order to minimize request collisions, it calculates t_6 as a random offset based on the Data Backoff Start value in the most recent Map (see Section 6.4.4, also the multicast SID definitions in Appendix, Section A.2).

At t_4 , the CM transmits a request for as many mini-slots as needed to accommodate the PDU. Time t_4 is chosen based on the ranging offset (see Section 6.3.2.5) so that the request will arrive at the CMTS at t_6 .

At t_6 , the CMTS receives the request and schedules it for service in the next map. (The choice of which requests to grant will vary with the class of service requested, any competing requests, and the algorithm used by the CMTS.)

At t_7 , the CMTS transmits a map whose effective starting time is t_9 . Within this map, a data grant for the CM will start at t_{11} .

At t_8 , the CM receives the map and scans for its data grant.

At t_{10} , the CM transmits its data PDU so that it will arrive at the CMTS at t_{11} . Time t_{10} is calculated from the ranging offset as in step 3.

Steps 1 and 2 need not contribute to access latency if CMs routinely maintain a list of request opportunities.

At Step 3, the request may collide with requests from other CMs and be lost. The CMTS does not directly detect the collision. The CM determines that a collision (or other reception failure) occurred when the next map fails to include acknowledgment of the request. The CM MUST then perform a back-off algorithm and retry.

At Step 4, the CMTS scheduler MAY fail to accommodate the request within the next map. If so, it MUST reply with a zero-length grant in that map or discard the request by giving no grant at all. It MUST continue to report this zero-length grant in all succeeding maps until the request can be granted or is discarded. This MUST signal to the CM that the request is still pending. So long as the CM is receiving a zero-length grant, it MUST NOT issue new requests for that service queue.

6.4.4 Contention Resolution

The CMTS controls assignments on the upstream channel through the MAP and determines which mini-slots are subject to collisions. The CMTS MAY allow collisions on either Requests or Data PDUs.

The mandatory method of contention resolution which MUST be supported is based on a truncated binary exponential back-off, with the initial back-off window and the maximum back-off window controlled by the CMTS. The values are specified as part of the Bandwidth Allocation Map (MAP) MAC message and represent a power-of-two value. For example, a value of 4 indicates a window between 0 and 15; a value of 10 indicates a window between 0 and 1023.

When a CM has information to send and wants to enter the contention resolution process, it sets its internal back-off window equal to the Data Backoff Start defined in the MAP currently in effect.

The CM MUST randomly select a number within its back-off window. This random value indicates the number of contention transmit opportunities which the CM MUST defer before transmitting. A CM MUST only consider contention transmit opportunities for which this transmission would have been eligible. These are defined by either Request IEs or Request/Data IEs in the MAP.

Note: Each IE can represent multiple transmission opportunities.

As an example, consider a CM whose initial back-off window is 0 to 15 and it randomly selects the number 11. The CM must defer a total of 11 contention transmission opportunities. If the first available Request IE is for 6 requests, the CM does not use this and has 5 more opportunities to defer. If the next Request IE is for 2 requests, the CM has 3 more to defer. If the third Request IE is for 8 requests, the CM transmits on the fourth request, after deferring for 3 more opportunities.

After a contention transmission, the CM waits for a Data Grant (Data Grant Pending) or Acknowledgment in a subsequent MAP. Once either is received, the contention resolution is complete. The CM determines that the contention transmission was lost when it finds a MAP without a Data Grant (Data Grant Pending) or Acknowledgment for it and with an Ack time more recent than the time of transmission. The CM MUST now increase its back-off window by a factor of two, as long as it is less than the maximum back-off window. The CM MUST randomly select a number within its new back-off window and repeat the deferring process described above.

This re-try process continues until the maximum number of retries (16) has been reached, at which time the PDU MUST be discarded.

Note: The maximum number of retries is independent of the initial and maximum back-off windows that are defined by the CMTS.

If the CM receives a unicast Request or Data Grant at any time while deferring for this SID, it MUST stop the contention resolution process and use the explicit transmit opportunity.

The CMTS has much flexibility in controlling the contention resolution. At one extreme, the CMTS MAY choose to set up the Data Backoff Start and End to emulate an Ethernet-style back-off with its associated simplicity and distributed nature, but also its fairness and efficiency issues. This would be done by setting initial = 0 and max = 10 in the MAP. At the other end, the CMTS MAY make the Data Backoff Start and End identical and frequently update these values in the MAP so all cable modems are using the same, and hopefully optimal, back-off window.

A Transmit Opportunity is defined as any mini-slot in which a CM may be allowed to start a transmission. Transmit Opportunities typically apply to contention opportunities and are used to calculate the proper amount to defer in the contention resolution process.

The number of Transmit Opportunities associated with a particular IE in a MAP is dependent on the total size of the region as well as the allowable size of an individual transmission. As an example, assume a REQ IE defines a region of 12 mini-slots. If the UCD defines a REQ Burst Size that fits into a single mini-slot, there are 12 Transmit Opportunities associated with this REQ IE, i.e., one for each mini-slot. If the UCD defines a REQ that fits in two mini-slots, there are six Transmit Opportunities and a REQ can start on every other mini-slot.

As another example, assume a REQ/Data IE that defines a 24 mini-slot region. If it is sent with a SID of 0x3FF4 (see Appendix A), a CM can potentially start a transmit on every fourth mini-slot; so this IE contains a total of six Transmit Opportunities (TX OP). Similarly, a SID of 0x3FF6 implies four TX OPs; 0x3FF8 implies three TX OPs; and 0x3FFC implies two TX OPs.

For an Initial Maintenance IE, a CM MUST start its transmission in the first mini-slot of the region; therefore it has a single Transmit Opportunity. The remainder of the region is used to compensate for the round trip delays since the CM has not yet been ranged.

Station Maintenance, Short and Long Data Grant IEs are specified for unicast and thus are NOT typically associated with contention Transmit Opportunities. They represent a single dedicated, or reservation based, Transmit Opportunity.

In summary:

Table 6-18. Transmit Opportunity

Interval	SID Type	Transmit Opportunity
Request	Broadcast	# minislots required for a Request
Request	Multicast	# minislots required for a Request
Request/Data	Broadcast	Not allowed
Request/Data	Well-known Multicast	As defined by SID in Appendix A
Request/Data	Multicast	Vendor specific algorithms
Initial Maint.	Broadcast	Entire interval is a single tx opp.
Initial Maint.	Multicast	Entire interval is a single tx opp.

6.4.5 CM Behavior

The following rules govern the response a CM may make when processing maps:

1. A CM MUST first use any Grants assigned to it. Next, the CM MUST use any unicast REQ for it. Finally, the CM MUST use the next available broadcast/multicast REQ or REQ/Data IEs for which it is eligible.
2. Only one Request may be outstanding at a time for a particular Service ID.
3. If a CM has a Request pending, it MUST NOT use intervening contention intervals for that Service ID.

6.4.6 Support for Multiple Channels

Vendors MAY choose to offer various combinations of upstream and downstream channels within one MAC service access point. The upstream bandwidth allocation protocol allows for multiple upstream channels to be managed via one or many downstream channels.

If multiple upstream channels are associated with a single downstream channel, the CMTS MUST send one allocation map per upstream channel. The map’s channel identifier, taken with the Upstream Channel Descriptor Message (see Section 6.3.2.2), MUST specify to which channel each map applies. There is no requirement that the maps be synchronized across channels. Appendix G provides an example.

If multiple downstream channels are associated with a single upstream channel, the CMTS MUST ensure that the allocation map reaches all CMs. That is, if some CMs are attached to a particular downstream channel, the map MUST be transmitted on that channel. This MAY necessitate that multiple copies of the same map be transmitted. The slot reference in the map header MUST always relate to the SYNC reference on the downstream channel on which it is transmitted.

If multiple downstream channels are associated with multiple upstream channels, the CMTS MAY need to transmit multiple copies of multiple maps to ensure both that all upstream channels are mapped and that all CMs have received their needed maps.

6.4.7 Classes of Service

The specification provides a few class of service parameters that can enhance or control service. These parameters provide the minimum functionality that the system requires, but could be enhanced through other mechanisms to provide even greater functionality. This section describes some approaches to different traffic types. Appendix C outlines the specific Class of Service parameters that MUST be supported and provides additional details.

This section illustrates how the available mechanisms can be used to provide support for the service classes defined in [RFC-1633] “Integrated Services in the Internet Architecture: An Overview”.

[RFC-1633] divides applications into elastic applications which will always wait for data to arrive and inelastic applications in which the data must arrive within a certain time to be useful.

Within the inelastic category further sub divisions can be defined:

- delay-intolerant—the data must arrive within a perfectly reliable upper bound on delay
- delay-tolerant—the data must arrive within a fairly reliable but not perfectly reliable delay bound.

Within the elastic category the following application types can be distinguished:

- interactive burst
- interactive bulk

The service model should be able to support both types of inelastic application and to allow for lower delays for interactive elastic applications than for bulk elastic applications.

Inelastic Delay-Intolerant—The CMTS provides a Data Grant of fixed size to a configured Service ID once every N mini-slots. This Service ID MAY be assigned to all traffic for a CM, or it MAY only be used for this particular service within the CM.

Inelastic Delay-Tolerant—The CMTS periodically provides a unicast Request IE to a configured Service ID. It then grants the request based on the negotiated delay variation, bandwidth, and other considerations. The CM has guaranteed access in which to make requests, and the CMTS’s scheduling algorithm provides the negotiated service. As an alternative, the minimum data rate of the service negotiation MAY be provided in the same way that inelastic delay-intolerant traffic is handled.

Elastic Application Support—is provided by a contention/FIFO service strategy, in which CMs contend for request slots, and the CMTS services requests as they arrive. Service priorities can allow differential delays between interactive and bulk applications.

6.4.7.1 Resource-Sharing

In order to support multiple end systems sharing the same upstream and downstream links, it is necessary to provide resource-sharing mechanisms for the link bandwidth. The following are some examples of this:

Link-usage feedback is provided implicitly by contention and by the CMTS’s scheduling algorithm, so no explicit congestion notifications are needed.

Service Priorities MUST be implemented by applying different service criteria to different Service IDs. It is anticipated that a particular CM MAY have several Service IDs, each corresponding to a particular service class. The particular services offered MAY vary from vendor to vendor.

Contention that is limited to a service class MAY be accomplished with multicast Request IEs and Request/Data IEs. Creation of such multicast groups is vendor-specific.

6.5 Timing and Synchronization

One of the major challenges in designing a MAC protocol for a cable network is compensating for the large delays involved. These delays are an order of magnitude larger than the transmission burst time in the upstream.

To compensate for these delays, the cable modem MUST be able to time its transmissions precisely to arrive at the CMTS at the start of the assigned mini-slot.

To accomplish this, two pieces of information are needed by each cable modem:

- a global timing reference sent downstream from the CMTS to all cable modems.
- a timing offset, calculated during a ranging process, for each cable modem.

6.5.1 Global Timing Reference

The CMTS MUST create a global timing reference by transmitting the Time Synchronization (SYNC) MAC management message downstream at a nominal frequency. The message contains a timestamp that exactly identifies when the CMTS transmitted the message. Cable modems MUST then compare the actual time the message was received with the timestamp and adjust their local clock references accordingly.

The Transmission Convergence sublayer must operate closely with the MAC sublayer to provide an accurate timestamp for the SYNC message. As mentioned in the Ranging section below (Section 6.5.3), the model assumes that the timing delays through the remainder of the PHY layer MUST be relatively constant. Any variation in the PHY delays MUST be accounted for in the guard time of the PHY overhead.

It is intended that the nominal interval between SYNC messages be tens of milliseconds. This imposes very little downstream overhead while letting cable modems acquire their global timing synchronization quickly.

6.5.2 CM Channel Acquisition

Any cable modem MUST NOT use the upstream channel until it has successfully synchronized to the downstream.

First, the cable modem MUST establish PMD sublayer synchronization. This implies that it has locked onto the correct frequency, equalized the downstream channel, recovered any PMD sublayer framing and the FEC is operational (refer to Section 7.2.1). At this point, a valid bit stream is being sent to the transmission convergence sublayer. The transmission convergence sublayer performs its own synchronization (see Section 5). On detecting the well-known DOCSIS PID, along with a payload unit start indicator per [ITU-T H.222.0], it delivers the MAC frame to the MAC sublayer.

The MAC sublayer MUST now search for the Timing Synchronization (SYNC) MAC management messages. The cable modem achieves MAC synchronization once it has received at least two SYNC messages and has verified that its clock tolerances are within specified limits.

A cable modem remains in “SYNC” as long as it continues to successfully receive the SYNC messages. If the Lost SYNC Interval (see B) has elapsed without a valid SYNC message, a cable modem MUST NOT use the upstream and MUST try to re-establish synchronization again.

6.5.3 Ranging

Ranging is the process of acquiring the correct timing offset such that the cable modem’s transmissions are aligned to the correct mini-slot boundary. The timing delays through the PHY layer MUST be relatively constant. Any variation in the PHY delays MUST be accounted for in the guard time of the upstream PMD overhead.

First, a cable modem MUST synchronize to the downstream and learn the upstream channel characteristics through the Upstream Channel Descriptor MAC management message. At this point, the cable modem MUST scan the Bandwidth Allocation MAP message to find an Initial Maintenance Region. Refer to Section 6.4.1.1.4. The CMTS MUST make an Initial Maintenance region large enough to account for the variation in delays between any two CMs.

The cable modem MUST put together a Ranging Request message to be sent in an Initial Maintenance region. The SID field MUST be set to the non-initialized CM value (zero).

Ranging adjusts each CM’s timing offset such that it appears to be located right next to the CMTS. The CM MUST set its initial timing offset to the amount of internal fixed delay equivalent to putting this CM next to the CMTS. This amount includes delays introduced through a particular implementation, and MUST include the downstream PHY interleaving latency.

When the Initial Maintenance transmit opportunity occurs, the cable modem MUST send the Ranging Request message. Thus, the cable modem sends the message as if it was physically right at the CMTS.

Once the CMTS has successfully received the Ranging Request message, it MUST return a Ranging Response message addressed to the individual cable modem. Within the Ranging Response message MUST be a temporary SID assigned to this cable modem until it has completed the registration process. The message MUST also contain information on RF power level adjustment and offset frequency adjustment as well as any timing offset corrections.

The cable modem MUST now wait for an individual Station Maintenance region assigned to its temporary SID. It MUST now transmit a Ranging Request message at this time using the temporary SID along with any power level and timing offset corrections.

The CMTS MUST return another Ranging Response message to the cable modem with any additional fine tuning required. The ranging request/response steps MUST be repeated until the response contains a Ranging Successful notification. At this point, the cable modem MUST join normal data traffic in the upstream. See Section 7 for complete details on the entire initialization sequence. In particular, state machines and the applicability of retry counts and timer values for the ranging process are defined in Section 7.2.4.

Note: The burst type to use for any transmission is defined by the Interval Usage Code (IUC). Each IUC is mapped to a burst type in the UCD message.

6.5.4 Timing Units and Relationships

The SYNC message conveys a time reference that is measured in 6.25-microsecond ticks. Additional resolution of 6.25/64 microseconds is also present in the SYNC message to allow the CM to track the CMTS clock with a small phase offset. These units were chosen as the greatest-common-divisor of the upstream mini-slot time across various modulations and symbol rates. As this is decoupled from particular upstream channel characteristics, a single SYNC time reference may be used for all upstream channels associated with the downstream channel.

The bandwidth allocation MAP uses time units of “mini-slots.” A mini-slot represents the byte-time needed for transmission of a fixed number of bytes. The mini-slot is expected to represent 16 byte-times, although other values could be chosen. The size of the mini-slot, expressed as a multiple of the SYNC time reference, is carried in the Upstream Channel Descriptor. The example in Table 6-20 relates mini-slots to the SYNC time ticks:

Table 6-19. Example Relating Mini-Slots to Time Ticks

Parameter	Example Value
Time tick	6.25 microseconds
Bytes per mini-slot	16 (nominal, when using QPSK modulation)
Symbols/byte	4 (assuming QPSK)
Symbols/second	2,560,000
Mini-slots/second	40,000
Microseconds/mini-slot	25
Ticks/mini-slot	4

The reader is encouraged to try other symbol rates and modulations. Note that the symbols/byte is a characteristic of an individual burst transmission, not of the channel. A mini-slot in this instance could represent either 16 or 32 bytes, depending on the modulation choice.

A “mini-slot” is the unit of granularity for upstream transmission opportunities. There is no implication that any PDU can actually be transmitted in a single mini-slot.

The MAP counts mini-slots in a 32-bit counter that normally counts to $(2^{32} - 1)$ and then wraps back to zero. The least-significant bits (i.e., bit 0 to bit 25-M) of the mini-slot counter MUST match the most-significant bits (i.e., bit 6+M to bit 31) of the SYNC timestamp counter. That is, mini-slot N begins at timestamp reference $(N * T * 64)$, where $T = 2^M$ is the UCD multiplier that defines the mini-slot (i.e., the number of timeticks per minislot).

Note: The unused upper bits of the 32-bit mini-slot counter (i.e., bit 26-M to bit 31) are not needed by the CM and MAY be ignored.

Note that the constraint that the UCD multiplier be a power of two has the consequence that the number of bytes per mini-slot must also be a power of two.

6.6 Data Link Encryption Support

The procedures to support data link encryption are defined in [DOCSIS8]. The interaction between the MAC layer and the security system is limited to the items defined below.

6.6.1 MAC Messages

MAC Management Messages (Section 6.3) MUST NOT be encrypted.

6.6.2 Framing

The following rules MUST be followed when encryption is applied to a data PDU:

- Security specific header (Privacy EH element of [DOCSIS8]) MUST be in the extended header and MUST be the first EH element of the Extended Header field (EHDR).
- Encrypted data is carried as Data PDU to the Cable MAC transparently.

7. CABLE MODEM - CMTS INTERACTION

This section covers the key requirements for the interaction between a CM and a CMTS. The interaction can be broken down into five basic categories: initialization, authentication, configuration, authorization, and signaling.

7.1 CMTS Initialization

The mechanism utilized for CMTS initialization (local terminal, file download, SNMP, etc.) is described in [DOCSIS5]. It MUST meet the following criteria for system interoperability.

- The CMTS MUST be able to reboot and operate in a stand-alone mode using configuration data retained in non-volatile storage.
- If valid parameters are not available from non-volatile storage or via another mechanism such as the Spectrum Management System (see [SMS]), the CMTS MUST not generate any downstream messages (including SYNC). This will prevent CMs from transmitting.
- The CMTS MUST provide the information defined in Section 6 to CMs for each upstream channel.

7.2 Cable Modem Initialization

The procedure for initialization of a cable modem MUST be as shown in Figure 7-1. This figure shows the overall flow between the stages of initialization in a CM. This shows no error paths, and is simply to provide an overview of the process. The more detailed finite state machine representations of the individual sections (including error paths) are shown in the subsequent figures. Timeout values are defined in B.

The procedure for initializing a cable modem and for a CM to reinitialize its MAC can be divided into the following phases:

- Scanning and synchronization to downstream
- Obtain upstream parameters
- Ranging and automatic adjustments
- Establish IP connectivity
- Establish time of day
- Transfer operational parameteres
- Registration
- Baseline Privacy initialization, if CM is provisioned to run Baseline Privacy

Each CM contains the following information when shipped from the manufacturer:

- A unique IEEE 802 48-bit MAC address which is assigned during the manufacturing process. This is used to identify the modem to the various provisioning servers during initialization.
- Security information as defined in [DOCSIS8] used to authenticate the CM.

The SDL (Specification and Description Language) notation used in the following figures is shown in Figure 7-2 (refer to ITU-T Recommendation Z.100 [ITU-T Z.100]).




Figure 7-1. CM Initialization Overview




Figure 7-2. SDL Notation

7.2.1 Scanning and Synchronization to Downstream

On initialization or after signal loss, the cable modem **MUST** acquire a downstream channel. The CM **MUST** have non-volatile storage in which the last operational parameters are stored and **MUST** first try to re-acquire this downstream channel. If this fails, it **MUST** begin to continuously scan the 6-MHz channels of the downstream frequency band of operation until it finds a valid downstream signal.

A downstream signal is considered to be valid when the modem has achieved the following steps:

- synchronization of the QAM symbol timing
- synchronization of the FEC framing
- synchronization of the MPEG packetization
- recognition of SYNC downstream MAC messages

While scanning, it is desirable to give an indication to the user that the CM is doing so.

7.2.2 Obtain Upstream Parameters

Refer to Figure 7-3. After synchronization, the CM **MUST** wait for an upstream channel descriptor message (UCD) from the CMTS in order to retrieve a set of transmission parameters for a possible upstream channel. These messages are transmitted periodically from the CMTS for all available upstream channels and are addressed to the MAC broadcast address. The CM **MUST** determine whether it can use the upstream channel from the channel description parameters.

The CM **MUST** collect all UCDs which are different in their channel ID field to build a set of useable channel IDs. If no channel can be found after a suitable timeout period, the CM **MUST** continue scanning to find another downstream channel.

The CM **MUST** determine whether it can use the upstream channel from the channel description parameters. If the channel is not suitable, the CM **MUST** try the next channel ID until it finds a useable channel. If the channel is

suitable, the CM MUST extract the parameters for this upstream from the UCD. It MUST then wait for the next SYNC message¹⁰ and extract the upstream mini-slot timestamp from this message. The CM then MUST wait for a bandwidth allocation map for the selected channel. It MAY begin transmitting upstream in accordance with the MAC operation and the bandwidth allocation mechanism.

The CM MUST perform initial ranging at least once per Figure 7-6. If initial ranging is not successful, the next channel ID is selected, and the procedure restarted from UCD extraction. When there are no more channel IDs to try, the CM MUST continue scanning to find another downstream channel.

¹⁰ Alternatively, since the SYNC message applies to all upstream channels, the CM may have already acquired a time reference from previous SYNC messages. If so, it need not wait for a new SYNC.




Figure 7-3. Obtaining Upstream Parameters

7.2.3 Message Flows During Scanning and Upstream Parameter Acquisition

The CMTS MUST generate SYNC and UCD messages on the downstream at periodic intervals within the ranges defined in Section 6. These messages are addressed to all CMs. Refer to Figure 7-4.




Figure 7-4. Message Flows During Scanning and Upstream Parameter Acquisition

7.2.4 Ranging and Automatic Adjustments

The ranging and adjustment process is fully defined in Section 6 and in the following sections. The message sequence chart and the finite state machines on the following pages define the ranging and adjustment process which **MUST** be followed by compliant CMs and CMTSes. Refer to Figure 7-5 through Figure 7-8.

Note: MAPs are transmitted as described in Section 6.





Figure 7-5. Ranging and Automatic Adjustments Procedure

Note: The CMTS **MUST** allow the CM sufficient time to have processed the previous RNG-RSP (i.e., to modify the transmitter parameters) before sending the CM a specific ranging opportunity. This is defined as CM Ranging Response Time in B.



Note: Timeout T3 may occur because the RNG-REQs from multiple modems collided. To avoid these modems repeating the loop in lockstep, a random backoff is required. This is a backoff over the ranging window specified in the MAP. T3 timeouts can also occur during multi-channel operation. On a system with multiple upstream channels, the CM MUST attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.

Figure 7-6. Initial Ranging - CM





Figure 7-7. Initial Ranging - CM (continued)

Note: Ranging Request is within the tolerance of the CMTS.



1. Means ranging is within the tolerable limits of the CMTS.
2. RNG-REQ pending-till-complete was nonzero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CM's power level. If opportunities are offered prior to the pending-till-complete expiry, the "good-enough" test which follows receipt of a RNG-RSP MUST NOT judge the CM's transmit equalization until pending-till-complete expires.

Figure 7-8. Initial Ranging - CMTS

7.2.4.1 Ranging Parameter Adjustment

Adjustment of local parameters (e.g., transmit power) in a CM as a result of the receipt (or non-receipt) of an RNG-RSP is considered to be implementation-dependent with the following restrictions (refer to Section 6.2.7):

- All parameters MUST be within the approved range at all times
- Power adjustment MUST start from the minimum value unless a valid power is available from non-volatile storage, in which case this MUST be used as a starting point.
- Power adjustment MUST be capable of being reduced or increased by the specified amount in response to RNG-RSP messages.
- If, during initialization, power is adjusted to the maximum value it MUST wrap back to the minimum.
- For multi-channel support, the CM MUST attempt initial ranging on every suitable upstream channel before moving to the next available downstream channel.
- For multi-channel support, the CM MUST use the upstream channel ID of the range response as specified in Section 6.3.2.5 and in Appendix G.

7.2.4.2 Periodic Ranging

The CMTS MUST provide each CM a Periodic Ranging opportunity at least once every T4 seconds. The CMTS MUST send out Periodic Ranging opportunities at an interval sufficiently shorter than T4 that a MAP could be missed without the CM timing out. The size of this “subinterval” is CMTS dependent.

The CM MUST reinitialize its MAC after T4 seconds have elapsed without receiving a Periodic Ranging opportunity.

7.2.5 Establish IP Connectivity

At this point, the CM MUST invoke DHCP mechanisms [RFC-2131] in order to obtain an IP address and any other parameters needed to establish IP connectivity (see Appendix C). The DHCP response MUST contain the name of a file which contains further configuration parameters. Refer to Figure 7-9.




Figure 7-9. Establishing IP Connectivity

7.2.6 Establish Time of Day

The CM and CMTS need to have the current date and time. This is required for time-stamping logged events which can be reviewed by the management system. This need not be authenticated and need only be accurate to the nearest second.

The protocol by which the time of day **MUST** be retrieved is defined in [RFC-868]. Refer to Figure 7-10. The request and response **MUST** be transferred using UDP. The time retrieved from the server (UTC) **MUST** be combined with the time offset received from the DHCP response to create the current local time.

The DHCP server may offer a CM multiple Time of Day server IP addresses to attempt. The CM **MUST** attempt all Time of Day servers included in the DHCP offer until local time is established.

Successfully acquiring the time of Day is not mandatory for a successful registration, but it is necessary for ongoing operation. If a CM is unable to establish time of day before registration, it **MUST** log the failure, generate an alert to management facilities, then proceed to an operational state and retry periodically.

The specific timeout for Time of Day Requests is implementation dependent. However, for each server defined the CM **MUST NOT** exceed more than 3 Time of Day requests in any 5 minute period. At minimum, the CM **MUST** issue at least 1 Time of Day request per 5 minute period for each server specified until local time is established.]




Figure 7-10. Establishing Time of Day

7.2.7 Transfer Operational Parameters

After the DHCP and security association operations are successful, the modem **MUST** download the parameter file using TFTP, as shown in Figure 7-11. The TFTP configuration parameter server is specified by the “siaddr” field of the DHCP response. The CM **MUST** use an adaptive timeout for TFTP based on binary exponential backoff. Refer to [RFC-1123] and [RFC-2349].

See Appendix B, for the initial and final TFTP backoff values.

The parameter fields required in the DHCP response and the format and content of the configuration file **MUST** be as defined in Appendix C. Note that these fields are the minimum required for interoperability.

If a modem downloads a configuration file containing an upstream channel and/or downstream frequency different from what the modem is currently using, the modem **MUST NOT** send a Registration Request message to the CMTS. The modem **MUST** redo initial ranging using the configured upstream channel and/or downstream frequency per Section 6.3.2.5.3.

7.2.8 Registration

A CM **MUST** be authorized to forward traffic into the network once it is initialized, authenticated and configured. Refer to Figure 7-12.

The CM MUST forward its configured class of service and any other operational parameters in the configuration file (see Figure C.7) to the CMTS as part of a Registration Request. The CMTS MUST perform the following operations to confirm the CM authorization (refer to Figure 7-11):

- Calculate a MIC per C.6 and compare it to the CMTS MIC included in the Registration Request. If the MIC is invalid, the CMTS MUST respond with an Authorization Failure.
- If present and supported, check the TFTP Server Timestamp field. If the CMTS detects that the time is different from its local time by more than CM Configuration Processing Time (see Appendix B), the CMTS MUST indicate authentication failure in the REG-RSP. The CMTS SHOULD also make a log entry stating the CM MAC address from the message.
- If present and supported, check the TFTP Server Provisioned Modem Address field. If the Provisioned Modem Address does not match the requesting modem's actual address, the CMTS MUST indicate authentication failure in the REG-RSP. The CMTS SHOULD also make a log entry stating the CM MAC address from the message.
- Verify the availability of the class of service requested. If unable to provide the class of service, the CMTS MUST respond with a Class of Service Failure and the appropriate Service Not Available response code. (See C.7.16)
- Verify the availability of any Modem Capabilities requested. If unable or unwilling to provide the Modem Capability requested, the CMTS MUST turn that Modem Capability 'off' (see 6.3.2.7.1.1).
- Assign a Service ID for each class of service supported.
- Reply to the modem in a registration response.

Figure 7-12. shows the procedure that MUST be followed by the CM.




Figure 7-11. Registration - CMTS




Figure 7-12. Registration - CM

7.2.9 Baseline Privacy Initialization

Following registration, if the CM is provisioned to run Baseline Privacy, the CM MUST initialize Baseline Privacy operations, as described in [DOCSIS8].

7.2.10 Service IDs During CM Initialization

After completion of the Registration process (Section 7.2.8), the CM will have been assigned Service IDs (SIDs) to match its class of service provisioning. However, the CM must complete a number of protocol transactions prior to that time (e.g., Ranging, DHCP, etc.), and requires a temporary Service ID in order to complete those steps.

On reception of an Initial Ranging Request, the CMTS MUST allocate a temporary SID and assign it to the CM for initialization use. The CMTS MAY monitor use of this SID and restrict traffic to that needed for initialization. It MUST inform the CM of this assignment in the Ranging Response.

On receiving a Ranging Response addressed to it, the CM MUST use the assigned temporary SID for further initialization transmission requests until the Registration Response is received.

On receiving a Ranging Response instruction to move to a new downstream frequency and/or upstream channel ID, the CM MUST consider any previously assigned temporary SID to be deassigned, and must obtain a new temporary SID via initial ranging.

It is possible that the Ranging Response may be lost after transmission by the CMTS. The CM MUST recover by timing out and re-issuing its Initial Ranging Request. Since the CM is uniquely identified by the source MAC address in the Ranging Request, the CMTS MAY immediately re-use the temporary SID previously assigned. If the CMTS assigns a new temporary SID, it MUST make some provision for aging out the old SID that went unused (see Section 6.3.2.7).

When assigning class-of-service-provisioned SIDs on receiving a Registration Request, the CMTS may re-use the temporary SID, assigning it to one of the class of service classes requested. If so, it MUST continue to allow initialization messages on that SID, since the Registration Response could be lost in transit. If the CMTS assigns all-new SIDs for class-of-service provisioning, it MUST age out the temporary SID. The aging-out MUST allow sufficient time to complete the registration process in case the Registration Response is lost in transit.


7.2.11 Multiple-Channel Support

In the event that more than one downstream signal is present in the system, the CM MUST operate using the first valid downstream signal that it encounters when scanning. It will be instructed via the parameters in the configuration file (see Appendix C) to shift operation to different downstream and/or upstream frequencies if necessary.

Both upstream and downstream channels MUST be identified where required in MAC management messages using channel identifiers.

7.2.12 Remote RF Signal Level Adjustment

RF signal level adjustment at the CM is performed through a periodic maintenance function using the RNG-REQ and RNG-RSP MAC messages. This is similar to initial ranging and is shown in Figure 7-13 and Figure 7-14. On receiving a RNG-RSP, the CM MUST NOT transmit until the RF signal has been adjusted in accordance with the RNG-RSP and has stabilized (refer to Section 4).



Note 1: Means Ranging Request is within the tolerance limits of the CMTS for power and transmit equalization (if supported)

Note 2: RNG-REQ pending-till-complete was nonzero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CM's power level. If opportunities are offered prior to the pending-till-complete expiry, the "good-enough" test which follows receipt of a RNG-RSP MUST NOT judge the CM's transmit equalization until pending-till-complete expires.

Figure 7-13. Periodic Ranging - CMTS




Figure 7-14. Periodic Ranging - CM View

7.2.13 Changing Upstream Burst Parameters

Whenever the CMTS is to change any of the upstream burst characteristics, it must provide for an orderly transition from the old values to the new values by all CMs. Whenever the CMTS is to change any of the upstream burst values, it MUST:

- Announce the new values in an Upstream Channel Descriptor message. The Configuration Change Count field must be incremented to indicate that a value has changed.

After transmitting one or more UCD messages with the new value, the CMTS transmits a MAP message with a UCD Count matching the new Configuration Change Count. The first interval in the MAP MUST be a data grant of at least 1 millisecond to the null Service ID (zero). That is, the CMTS MUST allow one millisecond for cable modems to change their PMD sublayer parameters to match the new set. This millisecond is in addition to other MAP timing constraints (see Section 6.4.2).

- The CMTS MUST NOT transmit MAPs with the old UCD Count after transmitting the new UCD.

The CM MUST use the parameters from the UCD corresponding to the MAP's "UCD Count" for any transmissions it makes in response to that MAP. If the CM has, for any reason, not received the corresponding UCD, it cannot transmit during the interval described by that MAP.

7.2.14 Changing Upstream Channels

At any time after registration, the CMTS MAY direct the CM to change its upstream channel. This may be done for traffic balancing, noise avoidance, or any of a number of other reasons which are beyond the scope of this specification. Figure 7-15 shows the procedure that MUST be followed by the CMTS. Figure 7-16 shows the corresponding procedure at the CM.




Figure 7-15. Changing Upstream Channels: CMTS View

Note that if the CMTS retries the UCC-REQ, the CM may have already changed channels (if the UCC-RSP was lost in transit). Consequently, the CMTS MUST listen for the UCC-RSP on both the old and the new channels.




Figure 7-16. Changing Upstream Channels: CM View

The CM MUST successfully establish initial ranging on a new channel before using that channel. It MUST NOT perform re-registration, since its provisioning and MAC domain remain valid on the new channel. If the CM has previously established ranging on the new channel, and if that ranging on that channel is still current (T4 has not elapsed since the last successful ranging), the CM MAY use cached ranging information and may omit initial ranging.

7.2.15 Fault Detection and Recovery

Fault detection and recovery occurs at multiple levels.

- At the physical level, FEC is used to correct errors where possible. Refer to Section 4 for details.
- The MAC protocol protects against errors through the use of checksum fields across both the MAC Header and the data portions of the packet. Refer to Section 6 for details.
- All MAC management messages are protected with a CRC covering the entire message, as defined in Section 6. Any message with a bad CRC MUST be discarded by the receiver.

Table 7-1 shows the recovery process that MUST be taken following the loss of a specific type of MAC message.

Appendix I contains a list of error codes with more useful information as to the failure of the PHY and MAC layers. Refer to Section 6.2.7 for more information.

Table 7-1. Recovery Process on Loss of Specific MAC Messages

Message Name	Action Following Message Loss
SYNC	The CM can lose SYNC messages for a period of the loss SYNC interval (see B) before it has lost synchronization with the network. When this occurs, it follows the same procedures to reacquire connectivity as during initialization
UCD	During CM initialization the CM MUST receive a usable ¹ UCD before transmitting on the upstream. When in the 'Obtain Upstream Parameters' state of CM initialization process, if the CM doesn't receive a usable ² UCD within the T1 timeout period, the CM MUST NOT transmit on the upstream and MUST scan for another downstream channel. After receiving a usable UCD, whenever the CM receives an unusable UCD or a MAP with a UCD Count that doesn't match the Configuration Change Count of the last UCD received, the CM MUST NOT transmit on the upstream and MUST start the T1 timer. If the T1 timer expires under these circumstances, the CM MUST reset and reinitialize its MAC connection.
MAP	A CM MUST NOT transmit without a valid upstream bandwidth allocation. If a MAP is missed due to error, the CM MUST NOT transmit for the period covered by the MAP.
RNG-REQ RNG-RSP	If a CM fails to receive a valid ranging response within a defined timeout period after transmitting a request, the request MUST be retried a number of times (as defined in B). Failure to receive a valid ranging response after the requisite number of attempts MUST cause the modem to reset and reinitialize its MAC connection.
REG-REQ REG-RSP	If a CM fails to receive a valid registration response within a defined timeout period after transmitting a request, the request will be retried a number of times (as defined in B). Failure to receive a valid registration response after the requisite number of attempts will cause the modem to reset and reinitialize its MAC connection.
UCC-REQ UCC-RSP	If a CMTS fails to receive a valid upstream channel change response within a defined timeout period after transmitting a request, the request MUST be retried a number of times (as defined in B). Failure to receive a valid response after the requisite number of attempts MUST cause the CMTS to consider the CM as unreachable.

¹ A usable UCD is one that contains legal profiles that the modem can understand. The CM MAY also require that the UCD Count of the MAPs received match the Configuration Change Count field of the last received UCD before it considers the UCD as usable.

Messages at the network layer and above are considered to be data packets by the MAC Sublayer. These are protected by the CRC field of the data packet and any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

7.2.16 Prevention of Unauthorized Transmissions

A CM SHOULD include a means for terminating RF transmission if it detects that its own carrier has been on continuously for longer than the longest possible valid transmission.

8 SUPPORTING FUTURE NEW CABLE MODEM CAPABILITIES

8.1 Setting Up Communications on an Enhanced Basis

In the future, new types of CM or CMTS with enhanced characteristics may be introduced. Future-proofing is provided, in the protocols described herein, to permit these new types of CM or CMTS to set up communication on an enhanced basis.

Two methods are provided to accomplish this: one for use when the downstream channel supports upstream channels of varying capability and the other for the case where enhanced downstream channels are available.

8.1.1 Upstream Enhanced / Downstream Standard

The procedure for CMs with enhanced capabilities will be as follows.

- a) The enhanced CM acquires a standard downstream CMTS signal.
- b) The CM receives and interprets upstream channel descriptor (UCD) messages forwarded from the CMTS until it finds one for a channel with the enhanced characteristics which it wishes to use. It joins the upstream transmission stream of this channel which has been assigned to enhanced CMs in accordance with the information in the downstream CMTS signal.

8.1.2 Downstream Enhanced / Upstream Enhanced or Standard

The procedure for CMs with enhanced capabilities will be as follows.

- a) The enhanced CM acquires a standard downstream CMTS signal.
- b) The CM receives and interprets upstream channel descriptor (UCD) messages forwarded from the CMTS until it finds one for a channel with the best match to the enhanced characteristics which it wishes to use. It joins the upstream transmission stream of this channel which has been assigned to enhanced CMs in accordance with the information in the downstream CMTS signal.
- c) The enhanced CM interacts with the provisioning server for the purposes of agreeing upon the operating frequencies, modulation, data rate and other characteristics for enhanced operation.
- d) The enhanced CM changes operating frequencies and other characteristics accordingly, if necessary, and commences enhanced operation on a different downstream channel if necessary under conditions that will not interfere with the standard CMs.
- e) The CM acquires the new downstream CMTS signal and waits on appropriate UCD on this new channel.

8.2 Downloading Cable Modem Operating Software

A CMTS SHOULD be capable of being remotely reprogrammed in the field via a software download via the network.

The cable modem device MUST be capable of being remotely reprogrammed in the field via a software download over the network. This software download capability MUST allow the functionality of the cable modem to be changed without requiring that cable system personnel physically revisit and reconfigure each unit. It is expected that this field programmability will be used to upgrade cable modem software to improve performance, accommodate new functions and features (such as enhanced class of service support), correct any design deficiencies discovered in the software, and to allow a migration path as the Data Over Cable Interface Specification evolves.

The CM MUST use an adaptive timeout for TFTP based on binary exponential backoff (Refer to [RFC-1123] and [RFC-2349]). See Appendix B for the initial and final TFTP backoff values.

The mechanism used for download MUST be TFTP file transfer. The transfer MUST be initiated in one of two ways:

- An SNMP manager requests the CM upgrade. The details of the SNMP initiated process are included in [DOCSIS9] and MUST be supported.
- If the Software Upgrade File Name in the CM's configuration file does not match the current software image of the CM, the CM MUST request the specified file via TFTP from the Software Server.

Note: The Software Server IP Address is a separate configuration file parameter and MUST be present for the CM to download a new software image via this mechanism.

The CM MUST write the new software image to non-volatile storage. Once the file transfer is complete, the CM MUST restart itself with the new code image.

If a CM is unable to complete the software upgrade process for whatever reason, it MUST remain capable of accepting new software downloads.

The following list is a collection of failure modes that may occur during the software upgrade process. Each condition includes a required CM response.

- If a CM exhausts the required number of TFTP retries by issuing a total of 16 consecutive unsuccessful read requests for the same file, the CM MUST fall back to the last known working image and proceed to an operational state.
- If during or following the TFTP download, the CM determines that the image is damaged or corrupt, the CM MUST reject the image and MAY re-attempt the download up to the maximum number of TFTP sequence retries (ref. Appendix B). After its last failed download attempt, the CM MUST fall back to the last known working image and proceed to an operational state.
- If a CM successfully downloads (or detects during download) an image that is not intended for the device, it MUST reject the file, fall back to the last known working image and proceed immediately to an operational state.

In addition, a CM MUST log any software upgrade process failure and MUST report it asynchronously to the network manager. All alert and log messages MUST include text that explicitly describes the specific condition resulting in failure.

Upon successful completion of a software upgrade, the CM MUST log the upgrade and MUST report it asynchronously to a network manager.

A CMTS MUST permit forwarding of software image downloads to a CM either before or after successful registration.

Following upgrade of the operational software, the CM MAY need to follow one of the procedures described above in order to change channels to use the enhanced functionality.

If the CM is to continue to operate in the same upstream and downstream channels as before the upgrade, it MUST be capable of inter-working with other CMs which MAY be running previous releases of software.

Where software has been upgraded to meet a new version of the specification, it is critical that it MUST inter-work with the previous version in order to allow a gradual transition of units on the network.

The periodic SYNC message transmitted on the downstream channel MUST indicate the protocol revision at which the channel is operating.

9 PROVISION FOR OTHER FUTURE CAPABILITIES

It is anticipated that cable modem networks will, in the future, support capabilities that cannot be adequately defined today. These capabilities may include:

- New physical-layer modulation encoding.
- Improvements to, or new configuration settings within, the defined physical-layer encoding.
- Differing traffic flows and classes of service (e.g., STM telephony).

It is the intent of this specification to provide for interoperability with future devices and networks to whatever extent is practical. The minimum level of interoperability is that future-capability modems and modems conforming to this specification are assigned to different frequency bands, and all modems can automatically scan to find a congenial frequency band.

9.1 Anticipated Physical-Layer Changes

Existing MAC signaling provides for optional transmitter equalization (see Section 6.3.2.5).

Other forms of upstream transmission manipulation, such as Tomlinson-Harashima precoding, may be developed in the future. Signaling to support such can be added as optional TLV-encodings for the Ranging Response message.

This configuration setting can be phased into existing networks without placing new requirements on existing devices.

When developing a new network, it may be necessary to know modem capabilities before coming to rely on a feature like this. The “Modem Capabilities” mask, exchanged as part of the CM-to-CMTS registration process (see Figure 6.3.2.7) is intended to provide this information.

9.1.1 Adding Upstream Channel and Burst Configuration Settings

In future, configuration settings may be provided for new upstream channel characteristics:

- Higher-symbol-rate signaling

and new upstream burst characteristics:

- Trellis-coded modulation (2 bits/symbol and 4 bits/symbol)
- Interleaving within a burst.

These are defined through new encodings of the Upstream Channel Descriptor. A CM which finds characteristics which it does not implement is required to either abstain from that burst type, or to find a different upstream channel (see Section 8.1.1). This is also controllable by administrative policy if enough commonality is present to complete the registration process.

As with transmission precoding, a modem-capabilities flag may be needed if the CMTS is to choose least-common-denominator capability.

9.1.1.1 Channel Burst Parameters for Advanced Modems

Configuration settings for channel burst parameters for advanced modems are given in Table 9-1.

Table 9-1. Channel Burst Parameters for Advanced Modems

Parameter	Configuration Settings
Modulation (additional configuration settings)	Trellis Coded Modulation available 1) 8PSK--2 bits/s (analogous to QPSK) 2) 32QAM--4 bits/s (analogous to 16QAM) 2 encoder configuration settings available for each.
Interleaving N rows by M columns transmitter fills columns	N = 0 to 255; 0 = no interleaving M = 1 to 256
Tomlinson-Harashima Precoding	(1) TH Precoding (2) Conventional Transmit FIR Equalization (3) None

It should be possible to program these capabilities separately to users on a given channel. For example, two users should be able to be commanded to operate at a given channel frequency and symbol rate, with one user having any or all of these features: 8PSK TCM, Interleaving, and TH Precoding; while the other user employs QPSK and none of the other features (i.e., this user is not an advanced cable modem).

9.1.2 Downstream Channel Improvements

Downstream channel improvements may require additional frequencies to implement for interoperability. The modem initialization process defined herein provides that, if the CM is unable to complete satisfactory exchanges with the CMTS, it will scan for a more suitable frequency (see Section 8.1.2).

9.2 New Network Service Requirements

The types of network service expected on a cable network are apt to change over the lifetime of equipment conforming to this specification. This specification anticipates use of ATM-style traffic parameters by giving the CMTS centralized control over bandwidth allocation and jitter. Future networks may include classes of data other than those explicitly provided (802-like and ATM). These may be implemented by using the Reserved code point in the MAC FC field. Because this specification does not require a particular bandwidth allocation algorithm, future algorithms may be developed which take into account policies and traffic types that are not yet well-understood.

9.2.1 Multicast Service IDs

Multicast Service IDs provide extensibility to the interval usage codes that are defined herein in the upstream bandwidth allocation map. The multicast ID reflects, not just group membership, but also the access rules that apply to whatever interval is assigned to that ID. The following examples of Request/Data IEs illustrate some of the possibilities for use of a particular ID:

- The grant is for contention space for all high-priority (as defined locally) data PDUs from a select group of CMs.
- The grant is for ATM cells only.

It may be necessary to develop an extension to the MAC signaling protocol to distribute the definition of attributes associated with particular multicast Service IDs.

9.2.2 RSVP Support for Upstream Traffic

The Resource ReSerVation Protocol (RSVP) is a resource reservation setup protocol currently being standardized by the Internet Engineering Task Force. RSVP provides receiver-initiated setup of resource reservations for multicast and unicast data flows. This section serves to anticipate and guide the definition of new MAC management messages to support resource reservation for upstream traffic in the Data-over-Cable context.

RSVP assumes the implementation of two modules on each RSVP-capable node to forward data packets: the “packet classifier” and the “packet scheduler”. The packet classifier determines the route and class-of-service class for each packet, and sends the packet to the packet scheduler. The RSVP packet classifier uses a “filter spec” (which matches a particular source IP address and TCP/UDP port number) to classify and restrict traffic that consumes reservation resources. The packet scheduler makes packet forwarding decisions (e.g., queuing decisions) to achieve the promised class of service on the interface. The RSVP packet scheduler uses a “flow spec” (which identifies token bucket parameters, peak data rate, etc.) to identify the desired class of service.

In the context of RSVP for upstream traffic in the data-over-cable system, it is desirable for the CM to perform the “packet classifier” function; however the CMTS should perform most of the “packet scheduler” function. The support for this split of functions suggests the future definition of three new MAC management messages: “Dynamic Service Addition”, “Dynamic Service Deletion”, and “Dynamic Service Response.”

The Dynamic Service Addition message is periodically transmitted from the CMTS to the CM to announce the allocation of a new SID. The Dynamic Service Addition message contains the new SID value, and type/length/ value fields which can encode the RSVP filter specification and RSVP “cleanup timeout” interval (to support the RSVP “soft state” approach). The CM is expected to use the new SID exclusively for upstream traffic that matches the filter specification. The CM should assume that the new SID is refreshed by the receipt of another Dynamic Service Addition message within the cleanup timeout interval; otherwise, the SID is ignored by the CM at the conclusion of the interval.

The Dynamic Service Deletion message is transmitted from the CMTS to the CM to delete an unused SID immediately (to support the RSVP explicit “teardown” message). The Dynamic Service Response message is transmitted from the CM to the CMTS to acknowledge receipt of a Dynamic Service Addition or Dynamic Service Deletion message.

The interaction between RSVP “Path” and “Resv” messages, and the Dynamic Service Addition and Dynamic Service Response messages, is proposed to be as follows:

1. The data flow source-node (behind a CM) generates an RSVP Path message, and sends the message toward the data flow destination-node.
2. The CM forwards the upstream RSVP Path message to the destination-node without processing.
3. The CMTS intercepts the upstream RSVP Path message, stores “path state” from the message, updates the “previous hop address” in the message, and forwards the message.
4. The data flow destination-node receives the RSVP Path message, and replies with an RSVP Resv message to request a reservation of resources for the data flow from the source-node to itself. The RSVP Resv message is sent to the “previous hop” of the Path message - toward the CMTS.
5. The CMTS receives the downstream RSVP Resv message, and processes the message flow spec using its “admission control” and “policy control” modules (in cooperation with the CMTS upstream bandwidth scheduler). Note that the CMTS must account for available resources on the CMTS-NSI as well as the RFI. The rest of this section assumes that the reservation message is accepted by the CMTS.
6. The CMTS sends the “Dynamic Service Addition” MAC message to the CM. The message includes a new SID and the “filter spec” from the RSVP Resv message.
7. The CM receives the “Dynamic Service Addition” MAC message, stores the new SID and “filter spec”, and sends the “Dynamic Service Response” MAC message to the CMTS.

8. The CMTS receives the “Dynamic Service Response” MAC message, and forwards the RSVP Resv message to its “previous hop.”
9. The CM forwards the downstream RSVP Resv message to the source-node without processing.

9.3 PID Filtering Capability

This specification uses a single well-known PID for all data-over-cable traffic. CMs MAY use additional PIDs for differentiation of traffic types or to provide streams to individual CMs. PID assignments MAY be facilitated by the appropriate MAC control message extensions. As an example, this could facilitate services that use MPEG packet-level encryption. Any such services are beyond the scope of this version of the specification.

An additional modem capability configuration setting could be added in the Registration Request (REG-REQ) message to indicate the number of PIDs, in addition to the well-known PID, that the CM can filter. A “0” would indicate that the CM can only filter on the well-known PID.

An extension to the encodings in the Registration Response (REG-RSP) could be used to assign to a CM additional PIDs on which to filter.

Appendix A. Well-known Addresses

A.1 MAC Addresses

MAC addresses described here are defined using the Ethernet/[ISO8802-3] convention as bit-little-endian.

The following multicast address **MUST** be used to address the set of all CM MAC sublayers; for example, when transmitting Allocation Map PDUs.

01-E0-2F-00-00-01

The following multicast address **MUST** be used to address all CMTSes within the MAC-sublayer domain:

01-E0-2F-00-00-02

Note that in nearly all cases the unicast CMTS address is preferred. The address range

01-E0-2F-00-00-03 through 01-E0-2F-00-00-0F

is reserved for future definition. Frames addressed to any of these addresses **SHOULD NOT** be forwarded out of the MAC-sublayer domain.

A.2 MAC Service IDs

The following MAC Service IDs have assigned meanings. Those not included in this table are available for assignment, either by the CMTS or administratively.

0x0000 Addressed to no CM.

0x3FFF Addressed to all CMs.

0x3FF1-0x3FFE Addressed to all CMs. Available for small data PDUs, as well as requests (used only with request/data IEs). The last digit indicates the frame length and transmission opportunities as follows:

0x3FF1 Within the interval specified, a transmission may start at any mini-slot, and must fit within one mini-slot.

0x3FF2 Within the interval specified, a transmission may start at every other mini-slot, and must fit within two mini-slots (e.g., a station may start transmission on the first mini-slot within the interval, the third mini-slot, the fifth, etc.).

0x3FF3 Within the interval specified, a transmission may start at any third mini-slot, and must fit within three mini-slots (e.g., starts at first, fourth, seventh, etc.).

0x3FF4 Starts at first, fifth, ninth, etc.

0x3FFD Starts at first, fourteenth (14th), twenty-seventh (27th), etc.

0x3FFE Within the interval specified, a transmission may start at any 14th mini-slot, and must fit within 14 mini-slots.

A.3 MPEG PID

All DOCSIS data **MUST** be carried in MPEG-2 packets with the header PID field set to 0x1FFE.

Appendix B. Parameters and Constants

System	Name	Time Reference	Minimum Value	Default Value	Maximum Value
CMTS	Sync Interval	Nominal time between transmission of SYNC messages (ref 6.3.2.1)			200 msec
CMTS	UCD Interval	Time between transmission of UCD messages (ref. 6.3.2.2)			2 sec
CMTS	Max MAP Pending	The number of mini-slots that a CMTS is allowed to map into the future (ref. 6.3.2.3)			4096 mini-slot times
CMTS	Ranging Interval	Time between transmission of broadcast Ranging requests (ref. 6.3.2.4)			2 sec
CM	Lost Sync Interval	Time since last received Sync message before synchronization is considered lost			600 msec
CM	Contention Ranging Retries	Number of Retries on contention Ranging Requests (ref. 7.2.4)		16	
CM, CMTS	Invited Ranging Retries	Number of Retries on inviting Ranging Requests (ref. 7.2.4)		16	
CM	Request Retries	Number of retries on bandwidth allocation requests		16	
CM	Registration Request Retries	Number of retries on registration requests		3	
CM	Data Retries	Number of retries on immediate data transmission		16	
CMTS	CM MAP processing time	Time provided between arrival of the last bit of a MAP at a CM and effectiveness of that MAP (ref. 6.4.1)	200 μ s		
CMTS	CM Ranging Response processing time	Minimum time allowed for a CM following receipt of a ranging response before it is expected to reply to an invited ranging request	1 msec		
CMTS	CM Configuration	The maximum time allowed for a CM, following receipt of a configuration file, to send a Registration Request to a CMTS.	3 minutes		
CM	T1	Wait for UCD timeout			5 * UCD interval maximum value
CM	T2	Wait for broadcast ranging timeout			5 * ranging interval
CM	T3	Wait for ranging response	50 msec	200 msec	200 msec
CM	T4	Wait for unicast ranging opportunity. If the pending-till-complete field was used earlier by this modem, the value of that field must be added to this interval.	30 sec		35 sec
CMTS	T5	Wait for Upstream Channel Change response			2 sec
CM	T6	Wait for registration response			3 sec
CM, CMTS	Mini-slot size	Size of mini-slot for upstream transmission. Must be a power of 2 (in units of the Timebase Tick)	32 symbol times		
CM, CMTS	Timebase Tick	System timing unit		6.25 μ sec	
CM	TFTP Backoff Start	Initial value for TFTP backoff	1 sec		
CM	TFTP Backoff End	Last value for TFTP backoff	16sec		
CM	TFTP Retries	Number of Retries on TFTP request		16	
CM		Number of retries on entire TFTP downloads	1		3
CMTS	Short or Long Data Grant Size	Size of short or long data grant defined within a single Allocation Map Interval			255 mini-slots

Note that only the initialization intervals must be aligned. Once the CM is assigned its proper channel ID, its activities need only be aligned with other users of its fiber branch. Ordinary data transmission and requests for bandwidth may occur independently across the four upstream channels.

Appendix C. CM Configuration Interface Specification

C.1 DHCP Fields Used by the CMTS and CM

The following fields MUST be present in the DHCP request from the CM and MUST be set as described below:

- The hardware type (htype) MUST be set to 1 (Ethernet).
- The hardware length (hlen) MUST be set to 6.
- The client hardware address (chaddr) MUST be set to the 48-bit MAC address associated with the RF interface of the CM.
- The “client identifier” option MUST be included, with the hardware type set to 1, and the value set to the same 48-bit MAC address as the chaddr field.
- Option code 60 (vendor class identifier). The CM MAY include the string “docsis 1.0” in this field
- The “parameter request list” option MUST be included. The option codes that MUST be included in the list are:
 - Option code 1 (Subnet Mask)
 - Option code 2 (Time Offset)
 - Option code 3 (Router Option)
 - Option code 4 (Time Server Option)
 - Option code 7 (Log Server Option)

The following fields are expected in the DHCP response returned to the CM. The CM MUST configure itself based on the DHCP response:

- The IP address to be used by the CM (yiaddr).
- The IP address of the TFTP server for use in the next phase of the bootstrap process (siaddr).
- If the DHCP server is on a different network (requiring a relay agent), then the IP address of the relay agent (giaddr).

Note: This may differ from the IP address of the first hop router.

- The name of the CM configuration file to be read from the TFTP server by the CM (file).
- The subnet mask to be used by the CM (Subnet Mask, option 1).
- The time offset of the CM from Universal Coordinated Time (UTC) (Time Offset, option 2). This is used by the CM to calculate the local time for use in time-stamping error logs.
- A list of addresses of one or more routers to be used for forwarding CM-originated IP traffic (Router Option, option 3). The CM is not required to use more than one router IP address for forwarding.
- A list of [RFC-868] time-servers from which the current time may be obtained (Time Server Option, option 4).
- A list of SYSLOG servers to which logging information may be sent (Log Server Option, option 7); see [DOCSIS5].

To assist the DHCP server in differentiating a CM discovery request from a CPE-side LAN discovery request, a CMTS MUST implement the following:

- All CMTSes MUST support the DHCP relay agent information option [RFC-3046]. Specifically, the CMTS MUST include the 48-bit MAC address of the RF side interface of the CM generating or bridging the DHCP discovery request in the agent remote ID sub-option field before relaying the discovery to a DHCP server.

- If the CMTS is a router, it MUST use a giaddr field to differentiate between CM and CPE side station if they are provisioned to be in different IP subnets. CMTSes SHOULD also provide this functionality.

C.2 CM Binary Configuration File Format

The CM-specific configuration data MUST be contained in a file which is downloaded to the CM via TFTP. This is a binary file in the same format defined for DHCP vendor extension data [RFC-2132].

It MUST consist of a number of configuration settings (1 per parameter) each of the form

type: length: value

where type is a single-octet identifier which defines the parameter

length is a single octet containing the length of the value field in octets (not including type and length fields)

value is from one to 254 octets containing the specific value for the parameter

The configuration settings MUST follow each other directly in the file, which is a stream of octets (no record markers).

Configuration settings are divided into three types:

- Standard configuration settings which MUST be present
- Standard configuration settings which MAY be present
- Vendor-specific configuration settings.

CMs MUST be capable of processing all standard configuration settings. CMs MUST ignore any configuration setting present in the configuration file which it cannot interpret.

Authentication of the provisioning information is provided by two message integrity check (MIC) configuration settings, CM MIC and CMTS MIC.

- CM MIC is a digest which ensures that the data sent from the provisioning server were not modified en route. This is NOT an authenticated digest (it does not include any shared secret).
- CMTS MIC is a digest used to authenticate the provisioning server to the CMTS during registration. It is taken over a number of fields one of which is a shared secret between the CMTS and the provisioning server.

Use of the CM MIC allows the CMTS to authenticate the provisioning data without needing to receive the entire file.

Thus the file structure is of the form shown in Figure C-1:




Figure C-1. Binary Configuration File Format

Note: Not all configuration settings need to be present in a given file.

C.3 Configuration File Settings

The following configuration settings MUST be included in the configuration file and MUST be supported by all CMs.

- Network Access Configuration Setting

- Class of Service Configuration Setting
- End Configuration Setting

The following configuration settings MAY be included in the configuration file and if present MUST be supported by all CMs.

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Baseline Privacy Configuration Setting
- Software Upgrade Filename Configuration Setting
- SNMP Write-Access Control
- SNMP MIB Object
- Software Server IP Address
- CPE Ethernet MAC Address
- Maximum Number of CPEs
- Pad Configuration Setting

The following configurations MAY be included in the configuration file and if present, and applicable to this type of modem, MUST be supported.

- Telephone Settings Option

The following configuration settings MAY be included in the configuration file and if present MAY be supported by a CM.

- Vendor-Specific Configuration Settings

C.4 Configuration File Creation

The sequence of operations required to create the configuration file is as shown in Figure C-1 through Figure C-5.

1. Create the type/length/value entries for all the parameters required by the CM.




Figure C-2. Create TLV Entries for Parameters Required by the CM

2. Calculate the CM message integrity check (MIC) configuration setting as defined in Section C.5 and add to the file following the last parameter using code and length values defined for this field.




Figure C-3. Add CM MIC

3. Calculate the CMTS message integrity check (MIC) configuration setting as defined in Section C.6 and add to the file following the CM MIC using code and length values defined for this field.




Figure C-4. Add CMTS MIC

4. Add the end of data marker.




Figure C-5. Add End of Data Marker

C.5 CM MIC Calculation

The CM message integrity check configuration setting **MUST** be calculated by performing an MD5 digest over the bytes of the configuration setting fields. It is calculated over the bytes of these settings as they appear in the TFTPed image, without regard to TLV ordering or contents. There are two exceptions to this disregard of the contents of the TFTPed image:

- 1) The bytes of the CM MIC TLV itself are omitted from the calculation. This includes the type, length, and value fields.

- 2) The bytes of the CMTS MIC TLV are omitted from the calculation. This includes the type, length, and value fields.

On receipt of a configuration file, the CM MUST recompute the digest and compare it to the CM MIC configuration setting in the file. If the digests do not match, the configuration file MUST be discarded.

C.6 CMTS MIC Calculation

The CMTS message integrity check configuration setting MUST be calculated by performing an MD5 digest over the following configuration setting fields, when present in the configuration file, in the order shown:

- Downstream Frequency Configuration Setting
- Upstream Channel ID Configuration Setting
- Network Access Configuration Setting
- Class of Service Configuration Setting
- Baseline Privacy Configuration Setting
- Vendor-Specific Configuration Settings
- CM MIC Configuration Setting
- Maximum Number of CPEs
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address

The bulleted list specifies the order of operations when calculating the CMTS MIC over configuration setting Type fields. The CMTS MUST calculate the CMTS MIC over TLVs of the same Type in the order they were received. Within Type fields, the CMTS MUST calculate the CMTS MIC over the Subtypes in the order they were received. To allow for correct CMTS MIC calculation by the CMTS, the CM MUST NOT reorder configuration file TLVs of the same Type or Subtypes within any given Type in its Registration-Request message.

All configuration setting fields MUST be treated as if they were contiguous data when calculating the CM MIC.

The digest MUST be added to the configuration file as its own configuration setting field using the CMTS MIC Configuration Setting encoding.

The authentication string is a shared secret between the provisioning server (which creates the configuration files) and the CMTS. It allows the CMTS to authenticate the CM provisioning. The authentication string is to be used as the key for calculating the keyed CMTS MIC digest as stated in C.6.1.

The mechanism by which the shared secret is managed is up to the system operator.

On receipt of a configuration file, the CM MUST forward the CMTS MIC as part of the registration request (REG-REQ).

On receipt of a REG-REQ, the CMTS MUST recompute the digest over the included fields and the authentication string and compare it to the CMTS MIC configuration setting in the file. If the digests do not match, the registration request MUST be rejected by setting the authentication failure result in the registration response status field.

C.6.1 Digest Calculation

The CMTS MIC digest field MUST be calculated using HMAC-MD5 as defined in [RFC-2104].

C.7 Encodings

The following type/length/value encodings MUST be used in both the configuration file and in CM registration requests. All multi-octet quantities are in network-byte order, i.e., the octet containing the most-significant bits is the first transmitted on the wire.

The following configuration settings MUST be supported by all CMs which are compliant with this specification.

C.7.1 End-of-Data Marker

This is a special marker for end of data.

It has no length or value fields.

Type

255

C.7.2 Pad Configuration Setting

This has no length or value fields and is only used following the end of data marker to pad the file to an integral number of 32-bit words.

Type

0

C.7.3 Downstream Frequency Configuration Setting

The receive frequency to be used by the CM. It is an override for the channel selected during scanning. This is the center frequency of the downstream channel in Hz stored as a 32-bit binary number.

Type	Length	RX Frequency
------	--------	--------------

1	4	rx1 rx2 rx3 rx4
---	---	-----------------

Valid Range

The receive frequency MUST be a multiple of 62500 Hz.

C.7.4 Upstream Channel ID Configuration Setting

The upstream channel ID which the CM MUST use. The CM MUST listen on the defined downstream channel until an upstream channel description message with this ID is found. It is an override for the channel selected during initialization.

Type	Length	Value
------	--------	-------

2	1	channel ID
---	---	------------

C.7.5 Network Access Control Object

If the value field is a 1 this CM is allowed access to the network; if a 0 it is not.

Type	Length	On / Off
3	1	1 or 0

C.7.6 Class of Service Configuration Setting

This field defines the parameters associated with a class of service. It is somewhat complex in that it is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the particular class of service parameters for the class of service in question. Note that the type fields defined are only valid within the encapsulated class of service configuration setting string. A single class of service configuration setting is used to define the parameters for a single service class. Multiple class definitions use multiple class of service configuration setting sets.

Type	Length	Value
4	n	

C.7.6.1 Internal Class of Service Encodings

C.7.6.1.1 Class ID

The value of the field specifies the identifier for the class of service to which the encapsulated string applies.

Type	Length	Value
1	1	

Valid Range

The class ID MUST be in the range 1 to 16.

C.7.6.1.2 Maximum Downstream Rate Configuration Setting

For a single SID modem, the value of this field specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

For a multiple SID modem, the aggregate value of these fields specifies the maximum downstream rate in bits per second that the CMTS is permitted to forward to CPE unicast MAC addresses learned or configured as mapping to the registering modem.

This is the peak data rate for Packet PDU Data (including destination MAC address and the CRC) over a one-second interval. This does not include MAC packets addressed to broadcast or multicast MAC addresses. The CMTS MUST limit downstream forwarding to this rate. The CMTS MAY delay, rather than drop, over-limit packets.

Type	Length	Value
2	4	

Note: This is a limit, not a guarantee that this rate is available.

C.7.6.1.3 Maximum Upstream Rate Configuration Setting

The value of this field specifies the maximum upstream rate in bits per second that the CM is permitted to forward to the RF Network.

This is the peak data rate for Packet PDU Data (including destination address and the CRC) over a one-second interval. The CM MUST limit all upstream forwarding (both contention and reservation-based), for the corresponding SID, to this rate. The CM MUST include Packet PDU Data packets addressed to broadcast or multicast addresses when calculating this rate.

The CM MUST enforce the maximum upstream rate. It SHOULD NOT discard upstream traffic simply because it exceeds this rate.

The CMTS MUST enforce this limit on all upstream data transmissions, including data sent in contention. The CMTS SHOULD generate an alarm if a modem exceeds its allowable rate.

Type	Length	Value
3	4	

Note: The purpose of this parameter is for the CM to perform traffic shaping at the input to the RF network and for the CMTS to perform traffic policing to ensure that the CM does not exceed this limit

The CMTS could enforce this limit by any of the following methods:

- a) discarding over-limit requests.
- b) deferring (through zero-length grants) the grant until it is conforming to the allowed limit.
- c) discarding over-limit data packets.
- d) Reporting to a policy monitor (for example, using the alarm mechanism) that is capable of incapacitating errant CMs.

Note: This is a limit, not a guarantee that this rate is available.

C.7.6.1.4 Upstream Channel Priority Configuration Setting

The value of the field specifies the relative priority assigned to this service class for data transmission in the upstream channel. Higher numbers indicate higher priority.

Type	Length	Value
4	1	

Valid Range

0 -> 7

C.7.6.1.5 Guaranteed Minimum Upstream Channel Data Rate Configuration Setting

The value of the field specifies the data rate in bit/sec which will be guaranteed to this service class on the upstream channel.

Type Length Value
 5 4

C.7.6.1.6 Maximum Upstream Channel Transmit Burst Configuration Setting

The value of the field specifies the maximum transmit burst (in bytes) which this service class is allowed on the upstream channel. A value of zero means there is no limit.

Note: This value does not include any physical-layer overhead.

Type Length Value
 6 2

Table C-1. Sample Class of Service Encoding

Type	Length	Value (sub)type	Length	Value	
4	28	1	1	1	class of service configuration setting service class 1
		2	4	10,000,000	max. forward rate of 10 Mb/sec
		3	4	2,000,000	max. return rate of 2 Mb/sec
		4	1	5	return path priority of 5
		5	4	64000	min guaranteed 64 kb/sec
		6	2	1518	max. Tx burst of 1518 bytes
4	28	1	1	2	class of service configuration setting service class 2
		2	4	5,000,000	max. forward rate of 5 Mb/sec
		3	4	1,000,000	max. return rate of 1 Mb/sec
		4	1	3	return path priority of 3
		5	4	32000	min guaranteed 32 kb/sec
		6	2	1518	max. Tx burst of 1518 bytes

C.7.6.1.7 Class-of-Service Privacy Enable

This configuration setting enables/disables Baseline Privacy on a provisioned CoS. See [DOCSIS8].

Type Length Enable / Disable
 7 (= CoS_BP_ENABLE) 1 1 or 0

C.7.7 Modem Capabilities Configuration Setting

The value field describes the capabilities of a particular modem, i.e., those OPTIONAL features which the modem can support. It is composed from a number of encapsulated type/length/value fields. The encapsulated fields define the specific capabilities for the modem in question. Note that the type fields defined are only valid within the encapsulated capabilities configuration setting string.

Type	Length	Value
5	n	

The set of possible encapsulated fields is described below.

C.7.7.1 Concatenation Support

If the value field is a 1 this modem can support concatenation; if a 0 it can not.

Type	Length	On / Off
1	1	1 or 0

Table C-2. Sample Capability Encoding

Type	Length	Value (sub)type	Length	Value	
5		1	1	1	modem capability configuration setting concatenation supported

C.7.7.2 Transmit Equalizer Taps per Symbol

This field shows the maximal number of pre-equalizer taps per symbol supported by the CM.

All CMs which support pre-equalization MUST support symbol-spaced equalizer coefficients. CM support of 2 or 4 taps per symbol is optional. If this tuple is missing, it is implied that the CM only supports symbol spaced equalizer coefficients.

Type	Length	Value
5.10	1	1, 2 or 4

C.7.7.3 Number of Transmit Equalizer Taps

This field shows the number of equalizer taps that are supported by the CM.

All CMs which support pre-equalization MUST support an equalizer length of at least 8 symbols. CM support of up to 64 T-spaced, T/2-spaced or T/4-spaced taps is optional. If this tuple is missing, it is implied that the CM only supports an equalizer length of 8 taps.

Type	Length	Value
5.11	1	8 to 64

C.7.8 CM Message Integrity Check (MIC) Configuration Setting

The value field contains the CM message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
6	16	d1 d2..... d16

C.7.9 CMTS Message Integrity Check (MIC) Configuration Setting

The value field contains the CMTS message integrity check code. This is used to detect unauthorized modification or corruption of the configuration file.

Type	Length	Value
7	16	d1 d2..... d16

C.7.10 Vendor ID Configuration Setting

The value field contains the vendor identification specified by the three-byte vendor-specific Organization Unique Identifier of the CM MAC address. When the vendor ID is used outside the vendor specific information field (VSIF) in the configuration file, it is the Vendor ID of the sender.

Type	Length	Value
8	3	v1, v2, v3

C.7.11 Software Upgrade Filename

The filename of the software upgrade file for the CM. The filename is a fully qualified directory-path name. The file is expected to reside on a TFTP server identified in a configuration setting option defined in Appendix C.7.23. See Section 8.2.

Type	Length	Value
9	n	filename

C.7.12 SNMP Write-Access Control

This object makes it possible to disable SNMP “Set” access to individual MIB objects. Each instance of this object controls access to all of the writeable MIB objects whose Object ID (OID) prefix matches. This object may be repeated to disable access to any number of MIB objects.

type	length	value
10	n	OID prefix plus control flag

where n is the size of the ASN.1 Basic Encoding Rules [ISO8025] encoding of the OID prefix plus one byte for the control flag.

The control flag may take values:

- 0 - allow write-access
- 1 - disallow write-access

Any OID prefix may be used. The Null OID 0.0 may be used to control access to all MIB objects. (The OID 1.3.6.1 will have the same effect.)

When multiple instances of this object are present and overlap, the longest (most specific) prefix has precedence. Thus, one example might be

someTable	disallow write-access
someTable.1.3	allow write-access

This example disallows access to all objects in someTable except for someTable.1.3.

C.7.13 SNMP MIB Object

This object allows arbitrary SNMP MIB objects to be Set via the TFTP-Registration process.

Type	Length	Value
11	n	variable binding

where the value is an SNMP VarBind as defined in [RFC-1157]. The VarBind is encoded in ASN.1 Basic Encoding Rules, just as it would be if part of an SNMP Set request.

The cable modem MUST treat this object as if it were part of an SNMP Set Request with the following caveats:

- It MUST treat the request as fully authorized (it cannot refuse the request for lack of privilege).
- SNMP Write-Control provisions (see previous section) do not apply.
- No SNMP response is generated by the CM.

This object MAY be repeated with different VarBinds to “Set” a number of MIB objects. All such Sets MUST be treated as if simultaneous.

Each VarBind MUST be limited to 255 bytes.

C.7.14 Vendor-Specific Information

Vendor-specific information for cable modems, if present, MUST be encoded in the vendor specific information field (VSIF) (code 43) using the Vendor ID field (C.7.10) to specify which TLV tuples apply to which vendors products. The Vendor ID MUST be the first TLV embedded inside VSIF. If the first TLV inside VSIF is not a Vendor ID, the TLV MUST be discarded.

This configuration setting MAY appear multiple times. The same Vendor ID MAY appear multiple times. However, there MUST NOT be more than one Vendor ID TLV inside a single VSIF.

Type	Length	Value
43	n	per vendor definition

Example:

Configuration with vendor A specific fields and vendor B specific fields:

VSIF (43) + n (number of bytes inside this VSIF)
 8 (Vendor ID Type) + 3 (length field) + Vendor ID of Vendor A
 Vendor A Specific Type #1 + length of the field + Value #1
 Vendor A Specific Type #2 + length of the field + Value #2

VSIF (43) + m (number of bytes inside this VSIF)
 8 (Vendor ID Type) + 3 (length field) + Vendor ID of Vendor B
 Vendor B Specific Type + length of the field + Value

C.7.15 Modem IP Address

This object informs the CMTS of the provisioned IP address of the cable modem.

Type	Length	Value
12	4	IP Address

This object appears only in the Registration Request message.

This address plays no part in the protocols defined in this specification, but is included to assist with network management. If the ‘TFTP Server Provisioned Modem Access’ option is included in the Registration Request message, the ‘Modem IP Address’ MAY be excluded.

Note: The ‘TFTP Server Provisioned Modem Address’ and the ‘Modem IP Address’ contain the same information. The difference is that the ‘Modem IP Address’ is not included in the CMTS MIC Calculation. Having distinct parameters is necessary since the ‘TFTP Server Provisioned Modem Address’ is optional.

C.7.16 Service(s) Not Available Response

This configuration setting MUST be included in the Registration Response message if the CMTS is unable or unwilling to grant any of the requested classes of service that appeared in the Registration Request. Although the value applies only to the failed service class, the entire Registration Request MUST be considered to have failed (none of the class-of-service configuration settings are granted).

Type	Length	Value
13	3	Class ID, Type, Reason

Where

Class ID	is the class-of-service class from the request which is not available
Type	is the specific class-of-service object within the class which caused the request to be rejected

Reason	is the reason for the rejection from the following:
	reason-other(1)
	reason-unrecognized-configuration-setting(2)
	reason-temporary(3)
	reason-permanent(4)

The reason codes MUST be used in the following way.

- Reason-other(1) is used when none of the other reason codes apply.
- Reason-unrecognized-configuration setting(2) is used when a class-of-service type is not recognized or when its value is outside of the specified range.
- Reason-temporary(3) indicates that the current loading of service IDs or traffic policies at the CMTS prevents granting the request, but that the request might succeed at another time.
- Reason-permanent(4) indicates that, for policy, configuration, or CMTS capabilities reasons, the request would never be granted unless the CMTS were manually reconfigured or replaced.

C.7.17 CPE Ethernet MAC Address

This object configures the CM with the Ethernet MAC address of a CPE device (see Section 3.1.2.3.1). This object may be repeated to configure any number of CPE device addresses.

Type	Length	Value
14	6	Ethernet MAC Address of CPE

This object appears only in the configuration file.

C.7.18 Telephone Settings Option

This configuration setting describes parameters which are specific to telephone return systems. It is composed from a number of encapsulated type/length/value fields. See [DOCSIS6].

Type	Length	Value
15 (= TRI_CFG01)	n	

C.7.19 Baseline Privacy Configuration Settings Option

This configuration setting describes parameters which are specific to Baseline Privacy. It is composed from a number of encapsulated type/length/value fields. See [DOCSIS8].

Type	Length	Value
17 (= BP_CFG)	n	

C.7.20 Maximum Number of CPEs

Table C-3 shows the maximum number of CPEs that can connect through a CM during a CM epoch. The CM epoch is (from Section 3.1.2.3.1) the time between startup and hard reset of the modem.

Note: This parameter should not be confused with the number of CPE addresses a CM may learn. A modem may learn Ethernet MAC addresses up to its maximum number of CPE addresses (from Section 3.1.2.3.1). The maximum number of CPEs that are granted access through the modem is governed by this configuration setting.

Table C-3. Maximum Number of CPEs

Type	Length	Description
18	1	The maximum number of CPEs that may be granted access through the CM.

If present, the value MUST be positive and non-zero. The non-existence of this option means the default value of 1.

Note: This is a limit on the maximum number of CPEs a CM will grant access to. Hardware limitations of a given modem implementation may require the modem to lower value.

C.7.21 TFTP Server Timestamp

Table C-4 shows the sending time of the configuration file in seconds. The definition of time is as in [RFC-868]

Table C-4. TFTP Server Timestamp

Type	Length	Description
19	4	Number of seconds since 00:00 1 Jan. 1900

C.7.22 TFTP Server Provisioned Modem Address

The IP Address of the modem requesting the configuration file.

Table C-5. TFTP Server Provisioned Modem Address

Type	Length	Description
20	4	IP Address of the CM

C.7.23 Software Upgrade TFTP Server

The IP address of the TFTP server, on which the software upgrade file for the CM resides. See Section 8.2 and Appendix C.7.11

Type	Length	Value
21	4	ip1,ip2,ip3,ip4

C.8 Summary

Table C-6. Default Configuration Values

Type	Length	Description	If Not Present
1	1	Downstream Frequency	The locked receive frequency should be used.
2	1	Upstream Channel ID	2-way RFI CM: default=current upstream channel Telephony Return CM: parameter is mandatory and must be set to 0 (zero).

3	1	(0/1) Network Access Control Object	(Mandatory)
4-X	n	Class of Service Encodings	2-way RFI CM: (Mandatory) Telephony Return CM: Default = Best Effort Service Note: telephony upstream is independent of CoS
4-1	1	[1.16] Class ID	(Mandatory)
4-2	4	Maximum Downstream Data Rate	No maximum downstream data rate
4-3	4	Maximum Upstream Data Rate	No maximum upstream data rate
4-4	1	[0.7] Upstream Channel Priority	0 (Lowest Priority)
4-5	4	Guaranteed Minimum Upstream Data Rate	0 (No guaranteed minimum upstream data rate)
4-6	2	Maximum Upstream Channel Burst	0 (No administrative limit on Maximum Upstream Channel Burst, but still bound by Upstream Channel Descriptor limits)
4-7	1	(0/1) C-o-S Privacy Enabled	0 (Privacy not enabled)
6	16	CM MIC	(Mandatory)
7	16	CMTS MIC	(Mandatory)
9	n	Software Upgrade Filename	Software upgrade is not needed
10	n	(0/1) SNMP Write Access Control	0 (No SNMP Write Access Control)
11	n	SNMP MIB Object	No SNMP objects set
15-X	n	Service Provider	(Refer to [DOCSIS6])
15-2	n	Service Providers	(Refer to [DOCSIS6])
15-3	n	Phone Number (1)	(Refer to [DOCSIS6])
15-4	n	Phone Number (2)	(Refer to [DOCSIS6])
15-5	n	Phone Number (3)	(Refer to [DOCSIS6])
15-6	1	Retry Connection	(Refer to [DOCSIS6])
15-7	n	Username	(Refer to [DOCSIS6])
15-8	n	Password	(Refer to [DOCSIS6])
15-9	1	DHCP Authenticate	(Refer to [DOCSIS6])
15-10	4	DHCP Server IP Address	(Refer to [DOCSIS6])
15-11	n	RADIUS Realm	(Refer to [DOCSIS6])
15-12	1	PPP Authentication	(Refer to [DOCSIS6])
15-13	4	Demand Dial Timer	(Refer to [DOCSIS6])
17	n	Baseline Privacy Configuration Settings	
17-1	4	Authorized Wait Timeout	(Refer to [DOCSIS8])
17-2	4	Re-authorize Wait Timeout	(Refer to [DOCSIS8])
17-3	4	Authorization Grace Time	(Refer to [DOCSIS8])
17-4	4	Operational Wait Timeout	(Refer to [DOCSIS8])
17-5	4	Re-key Wait Timeout	(Refer to [DOCSIS8])

17-6	4	TEK Grace Time	(Refer to [DOCSIS8])
17-7	4	Authorize Reject Wait Timeout	(Refer to [DOCSIS8])
18	1	Maximum # of CPEs	1
19	4	TFTP Server Timestamp	(Optional)
20	4	TFTP Server Provisioned Modem Address	(Optional)
21	4	Software Upgrade TFTP Server	IP Address of the server used to download this config file.
0	N/A	Pad	Not Applicable
255	N/A	End-of-Data	(Mandatory)

Appendix D. MAC Sublayer Service Definition

The MAC sublayer will provide the following services, consistent with [ISO/IEC15802-1]. This is an internal interface within the CM and CMTS and is provided for reference purposes only.

D.1 Service at the CM

The following service primitives are provided by the MAC sublayer to the higher-layer protocol entity. These represent an abstraction of the service provided and do not imply a particular implementation.

MAC_CM_802_DATA.request
 MAC_CM_DIX_DATA.request
 MAC_CM_ATM_DATA.request
 MAC_CM_802_DATA.indication
 MAC_CM_DIX_DATA.indication
 MAC_CM_ATM_DATA.indication
 MAC_CM_DATA.acknowledgment

D.2 MAC_CM_802_DATA.request

Issued by the higher-layer to request transmission from the CM on an upstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- service_ID
- contention_and_acknowledgment_constraints - specifies whether or not this request MAY be satisfied in a contention interval. Ordinarily, the CM will request that contention data be acknowledged by the CMTS.
- destination_address
- source_address (OPTIONAL) - if not explicitly overwritten, the address at this MSAP is used.
- LLC_pdu
- padding (OPTIONAL) - MAY be used if the LLC PDU is less than 60 bytes and it is desired to maintain [ISO8802-3] transparency.
- frame_check_sequence (OPTIONAL) - MAY be supplied if [ISO8802-3] transparency is desired. Otherwise, a 32-bit CRC sum is calculated by the MAC sublayer.
- length

D.3 MAC_CM_DIX_DATA.request

Issued by the higher-layer to request transmission from the CM on an upstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- service_ID
- contention_and_acknowledgment_constraints - specifies whether or not this request MAY be satisfied in a contention interval. Ordinarily, the CM will request that contention data be acknowledged by the CMTS.
- destination_address
- source_address (OPTIONAL) - if not explicitly overwritten, the address at this MSAP is used.
- ethernet type
- ethernet_dix_pdu

- length

D.4 MAC_CM_ATM_DATA.request

Issued by the higher-layer to request transmission from the CM on an upstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- service_ID
- contention_and_acknowledgment_constraints - specifies whether or not this request MAY be satisfied in a contention interval. Ordinarily, the CM will request that contention data be acknowledged by the CMTS.
- one or more ATM cells. Cells need not be within the same virtual circuit or virtual path.
- length

D.5 MAC_CM_802_DATA.indication

Issued by the CM MAC to indicate reception of data from the downstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- destination_address
- source_address
- LLC_pdu
- padding (OPTIONAL) - MAY be present if the LLC PDU was less than 60 bytes and [ISO8802-3] transparency was desired.
- frame_check_sequence
- length

D.6 MAC_CM_DIX_DATA.indication

Issued by the CM MAC to indicate reception of data from the downstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- destination_address
- source_address
- ethernet type
- ethernet_dix_pdu
- frame_check_sequence
- length

D.7 MAC_CM_ATM_DATA.indication

Issued by the CM MAC to indicate reception of data on the downstream channel. Parameters:

- channel_ID - MAY be implicit if the device supports attachment to a single channel.
- service_ID
- one or more ATM cells. Cells need not be within the same virtual circuit or virtual path.

- length

D.8 MAC_CM_DATA.acknowledgment

Issued by the CM MAC to indicate reception of an acknowledgment on the downstream channel. (An acknowledgment is an information element in a MAP PDU (see Section 6.4.1.1). The CMTS MUST include this IE in response to an upstream data transmission that includes an acknowledgment request.)

Parameters:

- channel_ID - The downstream channel on which the acknowledgment was received. May be implicit if the device supports attachment to a single channel.
- service_ID
- length

Appendix E. Example Burst Profiles

E.1 Introduction

Table E-1 through Table E-4 contain example Channel Burst Profiles for various modulation format and symbol rate combinations. The column labeled Column #1 in Table E-1 through Table E-4 corresponds to the Request burst type. The other columns correspond to the Communication (or Data) burst type. Table E-5 contains example Channel Burst Profiles corresponding to Power-Up burst types, or Acquisition burst types (for use on a new channel—or simply for refinement of user-unique parameters).

A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type, and will be assigned by the CMTS for each burst.

A programmable preamble superstring, up to 1024 bits long, is part of the channel-wide profile or attributes, common to the all burst profiles on the channel (Section 6.3.2.2, Table 6-14), but with each burst profile able to specify the start location within this sequence of bits and the length of the preamble (Section 6.3.2.2, Table 6-15). The first bit of the Preamble Pattern is designated by the Preamble Value Offset as described in Table 6-15, Section 6.3.2.2. The first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-8), and is I1 in the first symbol of the burst (see Section 4.2.2.2). As an example, per Table 6-15, for Preamble Offset Value = 100, the 101st bit of the preamble superstring is the first bit into the symbol mapper, and the 102nd bit is the second bit into the mapper, and is mapped to Q1, and so. An example 1024-bit-long preamble superstring is given in Section E.2.

Table E-6 contains the frame formats for each of the symbol rates with QPSK modulation for the example Request burst and for three codeword lengths for the Communication bursts, with one codeword per burst. Additionally, frame formats are shown for each of the rates with two of the example codeword lengths with four codewords per burst. In each format example, the information rate of the burst is calculated and given in the table. For the Request burst, the 6 bytes of “data” are assumed to be the information, and the rest is overhead. In the Communication bursts, the preamble, spacing (guard time), FEC parity, and the example 6 bytes of MAC Header are assumed overhead for the purposes of calculating information rate.

Table E-7 is structured the same as Table E-6, but with the example formats for 16QAM modulation.

E.2 Example Preamble Sequence

The following is the example 1024-bit preamble sequence for Table E-1 through Table E-5:

Bits 1 through 128:

```
1100 1100 1111 0000 1111 1111 1100 0000 1111 0011 1111 0011 0011 0000 0000 1100
0011 0000 0011 1111 1111 1100 1100 1100 1111 0000 1111 0011 1111 0011 1100 1100
```

Bits 129 through 256:

```
0011 0000 1111 1100 0000 1100 1111 1111 0000 1100 1100 0000 1111 0000 0000 1100
0000 0000 1111 1111 1111 0011 0011 0011 1100 0011 1100 1111 1100 1111 0011 0000
```

Bits 257 through 384:

```
1100 0011 1111 0000 0011 0011 1111 1100 0011 0011 0000 0011 1100 0000 0011 0000
0000 1110 1101 0001 0001 1110 1110 0101 0010 0101 0010 0101 1110 1110 0010 1110
```

Bits 385 through 512:

0010 1110 1110 0010 0010 1110 1110 1110 1110 1110 0010 0010 0010 1110 1110 0010
1110 1110 1110 0010 1110 0010 1110 0010 0010 0010 0010 1110 0010 0010 1110 0010

Bits 513 through 640:

0010 0010 1110 1110 1110 1110 1110 1110 0010 1110 0010 1110 0010 1110 1110 0010
0010 1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010

Bits 641 through 768:

0010 1110 1110 1110 0010 0010 0010 1110 0010 1110 1110 1110 1110 0010 0010 1110
0010 1110 0010 0010 0010 1110 1110 0010 0010 0010 0010 1110 0010 0010 0010 0010

Bits 769 through 896:

0010 1110 1110 1110 1110 1110 1110 0010 1110 0010 1110 0010 1110 1110 0010 0010
1110 1110 0010 1110 1110 1110 0010 1110 1110 0010 1110 0010 0010 1110 0010 0010

Bits 897 through 1024:

1110 1110 1110 0010 0010 0010 1110 0010 1110 1110 1110 1110 0010 0010 1110 0010
1110 0010 0010 0010 1110 1110 0010 0010 0010 0010 1110 0010 0010 0010 0010 1110

E.3 Example Burst Profiles

**Table E-1. Example Channel Burst Parameter Values for QPSK
Operation at 160, 320, and 640 ksym/sec**

Parameter	Config. Settings	#1	#2	#3	#4	#5
Modulation	QPSK, 16QAM	QPSK	QPSK	QPSK	QPSK	QPSK
Diff Enc	On/Off	Off	Off	Off	Off	Off
Symbol Rate	5 configuration settings	160, 320 or 640 ksym/ sec	160, 320 or 640 ksym/ sec			
Preamble Length	0, 4-1024 bits	56 bits	64 bits	64 bits	64 bits	64 bits
Preamble Start Location	1024 configuration settings	15	7	7	7	7
Preamble Values	1024 programmable bits	***	***	***	***	***
FEC On/Off	On/Off	Off	On	On	On	On
FEC Codeword Information Bytes (k)	1 to 255	N/A	32	56	64	220
FEC Error Correction (T bytes)	0 to 10	N/A	4	7	8	10
Last Codeword Length	Fixed or Shortened	N/A	Fixed	Fixed	Fixed	Fixed
Scrambler On/Off	On/Off	On	On	On	On	On
Scrambler Seed	15 bits**	default	default	default	default	default
Burst Length mini-slots*	0 to 255	3	0	0	0	0

*A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type.

**15 bits in a 16-bit field

***Refer to Section E.2.

Table E-2. Example Channel Burst Parameter Values for QPSK Operation at 1.28 and 2.56 Msym/sec

Parameter	Config. Settings	#1	#2	#3	#4	#5
Modulation	QPSK, 16QAM	QPSK	QPSK	QPSK	QPSK	QPSK
Diff Enc	On/Off	Off	Off	Off	Off	Off
Symbol Rate	5 configuration settings	1.28 or 2.56 Msym/ s				
Preamble Length	0, 4-1024 bits	48 bits	96 bits	96 bits	96 bits	96 bits
Preamble Start Location	1024 configuration settings	19	125	125	125	125
Preamble Values	1024 programmable bits	***	***	***	***	***
FEC On/Off	On/Off	Off	On	On	On	On
FEC Codeword Information Bytes (k)	1 to 255	N/A	40	56	64	220
FEC Error Correction (T bytes)	0 to 10	N/A	4	4	4	10
Last Codeword Length	Fixed or Shortened	N/A	Fixed	Fixed	Fixed	Fixed
Scrambler On/Off	On/Off	On	On	On	On	On
Scrambler Seed	15 bits**	default	default	default	default	default
Burst Length mini-slots*	0 to 255	4	0	0	0	0

*A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type.

**15 bits in a 16-bit field

***Refer to Section E.2.

Table E-3. Example Channel Burst Parameter Values for 16QAM Operation at 160, 320, and 640 ksym/sec

Parameter	Config. Settings	#1	#2	#3	#4	#5
Modulation	QPSK, 16QAM	16QAM	16QAM	16QAM	16QAM	16QAM
Diff Enc	On/Off	Off	Off	Off	Off	Off
Symbol Rate	5 configuration settings	160, 320 or 640 ksym/ sec				
Preamble Length	0, 4-1024 bits	80 bits	128 bits	128 bits	128 bits	128 bits
Preamble Start Location	1024 configuration settings	429	385	385	385	385
Preamble Values	1024 programmable bits	***	***	***	***	***
FEC On/Off	On/Off	Off	On	On	On	On
FEC Codeword Information Bytes (k)	1 to 255	N/A	32	56	64	220
FEC Error Correction (T bytes)	0 to 10	N/A	4	7	8	10
Last Codeword Length	Fixed or Shortened	N/A	Fixed	Fixed	Fixed	Fixed
Scrambler On/Off	On/Off	On	On	On	On	On
Scrambler Seed	15 bits**	default	default	default	default	default
Burst Length mini-slots*	0 to 255	2	0	0	0	0

*A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type.

**15 bits in a 16-bit field

***Refer to Section E.2.

**Table E-4. Example Channel Burst Parameter Values for 16QAM Operation
at 1.28 and 2.56 Msym/sec**

Parameter	Config. Settings	#1	#2	#3	#4	#5
Modulation	QPSK, 16QAM	16QAM	16QAM	16QAM	16QAM	16QAM
Diff Enc	On/Off	Off	Off	Off	Off	Off
Symbol Rate	5 configuration settings	1.28 or 2.56 Msym/ s				
Preamble Length	0, 4-1024 bits	144 bits	192 bits	192 bits	192 bits	192 bits
Preamble Start Location	1024 configuration settings	709	621	621	621	621
Preamble Values	1024 programmable bits	***	***	***	***	***
FEC On/Off	On/Off	Off	On	On	On	On
FEC Codeword Information Bytes (k)	1 to 255	N/A	40	56	64	220
FEC Error Correction (T bytes)	0 to 10	N/A	4	4	4	10
Last Codeword Length	Fixed or Shortened	N/A	Fixed	Fixed	Fixed	Fixed
Scrambler On/Off	On/Off	On	On	On	On	On
Scrambler Seed	15 bits**	default	default	default	default	default
Burst Length mini-slots*	0 to 255	4	0	0	0	0

*A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type.

**15 bits in a 16-bit field

***Refer to Section E.2.

Table E-5. Example Channel Burst Parameter Values for Power-Up and Acquisition in a New Channel

Parameter	Config. Settings	#1	#2	#3	#4
Modulation	QPSK, 16QAM	QPSK	QPSK	16QAM	16QAM
Diff Enc	On/Off	Off	Off	Off	Off
Symbol Rate	5 configuration settings	160, 320 or 640 ksym/sec	1.28 or 2.56 Msym/s	160, 320 or 640 ksym/sec	1.28 or 2.56 Msym/s
Preamble Length	0, 4-1024 bits	1024 bits	1024 bits	1024 bits	1024 bits
Preamble Start Location	1024 configuration settings	1	1	1	1
Preamble Values	1024 programmable bits	***	***	***	***
FEC On/Off	On/Off	On	On	On	On
FEC Codeword Information Bytes (k)	1 to 255	60	60	60	60
FEC Error Correction (T bytes)	0 to 10	10	10	10	10
Last Codeword Length	Fixed or Shortened	Fixed	Fixed	Fixed	Fixed
Scrambler On/Off	On/Off	On	On	On	On
Scrambler Seed	15 bits**	default	default	default	default
Burst Length mini-slots*	0 to 255	42	53	21	27

*A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type.

**15 bits in a 16-bit field

***Refer to Section E.2.

Table E-6. Frame Format Examples with QPSK Operation

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
Request Burst --					
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	24 (6)	24 (6)	24 (6)	24 (6)	24 (6)
preamble symbols (bytes)	28 (7)	28 (7)	28 (7)	24 (6)	24 (6)
total symbols (bytes)	60 (15)	60 (15)	60 (15)	64 (16)	64 (16)
total burst duration (mini-slots)	3	3	3	4	4
total burst duration (microseconds)	375	187.5	93.75	50	25
information rate (6 bytes per burst)	128 kb/s	256 kb/s	512 kb/s	960 kb/s	1.92 Mb/s
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	4	4	4	4	4
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	128 (32)	128 (32)	128 (32)	160 (40)	160 (40)
parity symbols (bytes)	32 (8)	32 (8)	32 (8)	32 (8)	32 (8)
preamble symbols (bytes)	32 (8)	32 (8)	32 (8)	48 (12)	48 (12)
total symbols (bytes)	200 (50)	200 (50)	200 (50)	256 (64)	256 (64)
total burst duration (mini-slots)	2+8=10	2+8=10	2+8=10	4+12=16	4+12=16
total burst duration (microseconds)	1250	625	312.5	200	100
information rate (excluding MAC Header)	166.4 kb/sec	332.8 kb/sec	665.6 kb/sec	1.360 Mb/ sec	2.720 Mb/ sec
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	8	8	8	4	4
spacing symbols (bytes)i.e., (guard time symbols -1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	256 (64)	256 (64)	256 (64)	256 (64)	256 (64)
parity symbols (bytes)	64 (16)	64 (16)	64 (16)	32 (8)	32 (8)
preamble symbols (bytes)	32 (8)	32 (8)	32 (8)	48 (12)	48 (12)
total symbols (bytes)	360 (90)	360 (90)	360 (90)	352 (88)	352 (88)

*The numbers in the table are given for a single codeword, but more codewords can be added, with the same data and parity lengths as given in the table, to create longer bursts.

(continued)

Table E-6. Frame Format Examples with QPSK Operation (continued)

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
total burst duration (mini-slots)	2+16=18	2+16=18	2+16=18	4+18=22	4+18=22
total burst duration (microseconds)	2250	1125	562.5	275	137.5
information rate (excluding MAC Header)	206.2 kb/sec	412.4 kb/sec	824.9 kb/sec	1.687 Mb/ sec	3.375 Mb/ sec
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	10	10	10	10	10
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	880 (220)	880 (220)	880 (220)	880 (220)	880 (220)
parity symbols (bytes)	80 (20)	80 (20)	80 (20)	80 (20)	80 (20)
preamble symbols (bytes)	32 (8)	32 (8)	32 (8)	48 (12)	48 (12)
total symbols (bytes)	1000 (250)	1000 (250)	1000 (250)	1024 (256)	1024 (256)
total burst duration (mini-slots)	2+48=50	2+48=50	2+48=50	4+60=64	4+60=64
total burst duration (microseconds)	6250	3125	1562.5	800	400
information rate (excluding MAC Header)	273.9 kb/sec	547.8 kb/sec	1.096 Mb/ sec	2.140 Mb/ sec	4.280 Mb/ sec
Communication Burst --					
codewords/burst	4*	4*	4*	4*	4*
errors corrected per codeword	8	8	8	4	4
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	1024 (256)	1024 (256)	1024 (256)	1024 (256)	1024 (256)
parity symbols (bytes)	256 (64)	256 (64)	256 (64)	128 (32)	128 (32)
preamble symbols (bytes)	32 (8)	32 (8)	32 (8)	48 (12)	48 (12)
total symbols (bytes)	1320 (330)	1320 (330)	1320 (330)	1216 (304)	1216 (304)
total burst duration (mini-slots)	2+16*4=66	2+16*4=66	2+16*4=66	4+18*4=76	4+18*4=76
total burst duration (microseconds)	8250	4125	2062.5	950	475
information rate (excluding MAC Header)	242.4 kb/sec	484.8 kb/sec	969.7 kb/sec	2.105 Mb/ sec	4.211 Mb/ sec

*The numbers in the table are given for a single codeword, but more codewords can be added, with the same data and parity lengths as given in the table, to create longer bursts.

(continued)

Table E-6. Frame Format Examples with QPSK Operation (continued)

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
Communication Burst --					

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
codewords/burst	4*	4*	4*	4*	1*
errors corrected per codeword	10	10	10	10	10
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (2)	8 (2)	8 (2)	16 (4)	16 (4)
data symbols (bytes)	3520 (880)	3520 (880)	3520 (880)	3520 (880)	3520 (880)
parity symbols (bytes)	320 (80)	320 (80)	320 (80)	320 (80)	320 (80)
preamble symbols (bytes)	32 (8)	32 (8)	32 (8)	48 (12)	48 (12)
total symbols (bytes)	3880 (970)	3880 (970)	3880 (970)	3904 (976)	3904 (976)
total burst duration (mini-slots)	2+48*4=194	2+48*4=194	2+48*4=194	4+60*4=244	4+60*4=244
total burst duration (microseconds)	24,250	12,125	6062.5	3050	1525
information rate (excluding MAC Header)	288.3 kb/sec	576.7 kb/sec	1.153 Mb/ sec	2.292 Mb/ sec	4.585 Mb/ sec

*The numbers in the table are given for four codewords per burst, but more or fewer codewords can be used, with the same data and parity lengths as given in the table.

Table E-7. Frame Format Examples with 16QAM Operation

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
Request Burst --					
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)
data symbols (bytes)	12 (6)	12 (6)	12 (6)	12 (6)	12 (6)
preamble symbols (bytes)	20 (10)	20 (10)	20 (10)	36 (18)	36 (18)
total symbols (bytes)	40 (20)	40 (20)	40 (20)	64 (32)	64 (32)
total burst duration (mini-slots)	2	2	2	4	4
total burst duration (microseconds)	250	125	62.5	50	25
information rate (6 bytes per burst)	192 kb/sec	384 kb/sec	768 kb/sec	960 kb/sec	1.920 Mb/ sec
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	4	4	4	4	4
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)
data symbols (bytes)	64 (32)	64 (32)	64 (32)	80 (40)	80 (40)
parity symbols (bytes)	16 (8)	16 (8)	16 (8)	16 (8)	16 (8)
preamble symbols (bytes)	32 (16)	32 (16)	32 (16)	48 (24)	48 (24)
total symbols (bytes)	120 (60)	120 (60)	120 (60)	160 (80)	160 (80)
total burst duration (mini-slots)	2+4=6	2+4=6	2+4=6	4+6=10	4+6=10
total burst duration (microseconds)	750	375	187.5	125	62.5
information rate (excluding MAC Header)	277.3 kb/sec	554.7 kb/sec	1.109 Mb/ sec	2.176 Mb/ sec	4.352 Mb/ sec

*The numbers in the table are given for a single codeword, but more codewords can be added, with the same data and parity lengths as given in the table, to create longer bursts.

(continued)

Table E-7. Frame Format Examples with 16QAM Operation (continued)

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	7	7	7	4	4
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)
data symbols (bytes)	128 (64)	128 (64)	128 (64)	128 (64)	128 (64)
parity symbols (bytes)	32 (16)	32 (16)	32 (16)	16 (8)	16 (8)
preamble symbols (bytes)	32 (16)	32 (16)	32 (16)	48 (24)	48 (24)
total symbols (bytes)	200 (100)	200 (100)	200 (100)	208 (104)	208 (104)
total burst duration (mini-slots)	2+8=10	2+8=10	2+8=10	4+9=13	4+9=13
total burst duration (microseconds)	1250	625	312.5	162.5	81.25
information rate (excluding MAC Header)	371.2 kb/sec	742.4 kb/sec	1.455 Mb/ sec	2.855 Mb/ sec	5.711 Mb/ sec
Communication Burst --					
codewords/burst	1*	1*	1*	1*	1*
errors corrected per codeword	10	10	10	10	10
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)
data symbols (bytes)	440 (220)	440 (220)	440 (220)	440 (220)	440 (220)
parity symbols (bytes)	40 (20)	40 (20)	40 (20)	40 (20)	40 (20)
preamble symbols (bytes)	32 (16)	32 (16)	32 (16)	48 (24)	48 (24)
total symbols (bytes)	520 (260)	520 (260)	520 (260)	544 (272)	544 (272)
total burst duration (mini-slots)	2+24=26	2+24=26	2+24=26	4+30=34	4+30=34
total burst duration (microseconds)	3250	1625	812.5	425	212.5
information rate (excluding MAC Header)	526.8 kb/sec	1.054 Mb/ sec	2.107 Mb/ sec	4.028 Mb/ sec	8.056 Mb/ sec

*The numbers in the table are given for a single codeword, but more codewords can be added, with the same data and parity lengths as given in the table, to create longer bursts.

(continued)

Table E-7. Frame Format Examples with 16QAM Operation (continued)

Parameter	160 ksym/ sec	320 ksym/ sec	640 ksym/ sec	1.28 Msym/ sec	2.56 Msym/ sec
Communication Burst --					
codewords/burst	4*	4*	4*	4*	4*
errors corrected per codeword	7	7	7	4	4
spacing symbols (bytes) i.e.,	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)

(guard time symbols - 1)					
data symbols (bytes)	512 (256)	512 (256)	512 (256)	512 (256)	512 (256)
parity symbols (bytes)	128 (64)	128 (64)	128 (64)	64 (32)	64 (32)
preamble symbols (bytes)	32 (16)	32 (16)	32 (16)	48 (24)	48 (24)
total symbols (bytes)	680 (340)	680 (340)	680 (340)	640 (320)	640 (320)
total burst duration (mini-slots)	$2+8*4=34$	$2+8*4=34$	$2+8*4=34$	$4+9*4=40$	$4+9*4=40$
total burst duration (microseconds)	4250	2125	1062.5	500	250
information rate (excluding MAC Header)	470.6 kb/sec	941.2 kb/sec	1.882 Mb/ sec	4.000 Mb/ sec	8.000 Mb/ sec
Communication Burst --					
codewords/burst	4*	4*	4*	4*	4*
errors corrected per codeword	10	10	10	10	10
spacing symbols (bytes) i.e., (guard time symbols - 1)	8 (4)	8 (4)	8 (4)	16 (8)	16 (8)
data symbols (bytes)	1760 (880)	1760 (880)	1760 (880)	1760 (880)	1760 (880)
parity symbols (bytes)	160 (80)	160 (80)	160 (80)	160 (80)	160 (80)
preamble symbols (bytes)	32 (16)	32 (16)	32 (16)	48 (24)	48 (24)
total symbols (bytes)	1960 (980)	1960 (980)	1960 (980)	1984 (992)	1984 (992)
total burst duration (mini-slots)	$2+24*4=98$	$2+24*4=98$	$2+24*4=98$	$4+30*4=124$	$4+30*4=124$
total burst duration (microseconds)	12,250	6125	3062.5	1550	775
information rate (excluding MAC Header)	570.8 kb/sec	1.142 Mb/ sec	2.283 Mb/ sec	4.511 Mb/ sec	9.022 Mb/ sec

*The numbers in the table are given for four codewords per burst, but more or fewer codewords can be used, with the same data and parity lengths as given in the table.

Appendix F. Upstream Modulation Rates

F.1 Introduction

The sources of ingress are not randomly distributed throughout the upstream cable network spectrum, but are grouped in frequency bands accordance with the International Radio Regulations, which is an international treaty subscribed to by most countries of the world. In some bands, higher powers are used, so ingress is more likely to be found there. The pattern of allocations can be readily recognized by examination of the upstream cable network band with a spectrum analyzer.

The upstream modulation rates and flexibility in selecting upstream carrier frequencies defined in this document permit operators to position carriers in those portions of the upstream cable network spectrum which are less prone to narrowband ingress.

F.2 Sources of Narrowband Ingress

Radio transmitters in the high-frequency (HF) portion of the radio spectrum (3 to 30 MHz) are the major sources of narrowband ingress. The transmitters that can affect upstream cable network paths may be local, domestic or in other countries.

F.2.1 HF Propagation

HF propagation is largely influenced by bending of radio waves in the ionosphere, which is a region between about 60 and 200 miles above the earth, so propagation distances can be very large. The bending is caused by ionization of free electrons by ultra-violet rays from the sun. There are several reflecting layers at various heights with differing characteristics, so the maximum usable frequency and distance reached depend on the time of day and solar activity.

Diurnal Variation - During local daytime, the higher frequencies are propagated over long distances and the lower frequencies are absorbed. During local night-time, the lower frequencies are propagated over distances which are not quite so large, and the higher frequencies experience diminished or no propagation.

Solar Cycle - Ionization is at its greatest during the peak of the 11-year sunspot cycle. At the largest sunspot peaks, strong distant signals can be heard across the HF band at most times of the day. At sunspot minima, long-distance propagation can be rare above about 15 MHz. Note that the cycle is currently at a minimum, so ingress is likely to affect cable systems on an increasing basis and at higher frequencies over the next 5 to 6 years.

F.2.2 Users of the HF Radio Spectrum

The main users of the HF radio spectrum are

- Broadcasting
- Maritime mobile
- Aeronautical mobile
- Aeronautical navigation
- Fixed (i.e., point-to-point)
- Standard frequency and time signals
- Amateur
- Citizens Band (CB)

The service of greatest concern for ingress is the Broadcasting Service. The Amateur and Citizens Band services are also of concern.

F.2.3 Broadcasting

The main ingress signals seen on typical cable networks are short-wave Broadcasting signals. The Broadcasting service is designed to provide very high field strengths in target markets so that the public can receive clear signals with simple receivers. Thus the service generally has these characteristics:

- Amplitude-modulated signals
- Very high transmitter powers
- High-gain, directive antenna arrays
- Simultaneous transmission on multiple frequencies in different parts of the HF spectrum, to combat HF propagation variability.
- Very large field strengths over a wide geographic area, affecting entire cable systems.

The frequencies allocated to the Broadcasting Service on a world-wide basis are shown in Table F-1.

Table F-1. Broadcasting Allocations Between 5 and 42 MHz

Frequency, kHz	
From	To
5,950	6,200
7,100	7,300
9,500	9,900
11,650	12,050
13,600	13,800
15,100	15,600
17,550	17,900
21,450	21,850
25,670	26,100

**Note: Although 7,100-7,300 kHz is not allocated to Broadcasting in ITU Region 2 (the Americas), strong signals reach the USA and Canada from the other ITU Regions.*

F.2.4 Amateur and CB

The Amateur Service is of concern, but for different reasons. Amateur transmitters operate at far lower powers than the Broadcasting Service, but are located in residential neighborhoods close to the cable plant, so they can represent potential localized sources of ingress in different bands across the HF spectrum. Radiation from the upstream cable plant may also cause interference into amateur receivers.

The Citizens Band Service operates at powers which are very much lower and are less likely to be sources of ingress. However, CB receivers may be capable of picking up interference from cable plant radiation.

The allocations in the USA and Canada for the Amateur and Citizens Band Services between 5 and 42 MHz are in Table F-2.

**Table F-2. Amateur and Citizens Band Allocations
Between 5 and 42 MHz**

Frequency, kHz	
From	To
7,000	7,300
10,100	10,150
14,000	14,350
18,068	18,168
21,000	21,450
24,890	24,990
26,960	27,240
28,000	29,700

F. 2.5 Other Services

The other services in the HF spectrum operate at powers which are lower than Broadcasting, and are usually not located in residential areas, so they are less likely to be of concern. In the low VHF region between 30 and 42 MHz, there may be a localized effect from high-power paging in the vicinity of 35 MHz.

F.3 Fitting Data Carriers Within the Ingress Gaps

The range of modulation rates and the flexibility in setting upstream carrier frequencies defined in this document permit operators to fit data-over-cable carriers in the gaps between the short-wave Broadcasting bands, to avoid this significant source of ingress. They also permit operators to avoid the Amateur and Citizens Band bands as well.

Table F-3 illustrates the quantity of data carriers that can be derived in the gaps between the Broadcasting allocations at the channel widths specified in this document, and the resulting utilization of the available spectrum. Table F-4 shows a similar illustration for the gaps between the Broadcasting, Amateur and CB allocations. A conservative upper limit of 40 MHz is used.

Table F.3. Data Carriers in the Gaps Between Broadcasting Bands

Non-Broadcast Spectrum, kHz			Data Carrier Width, kHz				
			200	400	800	1,600	3,200
From	To	Gap	Data Carriers Available				
5,000	5,950	950	4	2	1	0	0
6,200	7,100	900	4	2	1	0	0
7,300	9,500	2,200	11	5	2	1	0
9,900	11,650	1,750	8	4	2	1	0
12,050	13,600	1,550	7	3	1	0	0
13,800	15,100	1,300	6	3	1	0	0
15,600	17,550	1,950	9	4	2	1	0
17,900	21,450	3,550	17	8	4	2	1
21,850	25,670	3,820	19	9	4	2	1
26,100	40,000	13,900	69	34	17	8	4
Least Gap, kHz		900					
Total Carriers			154	74	35	15	6
Total Bandwidth, kHz		31,870	30,800	29,600	28,000	24,000	19,200
Utilization, %			97	93	88	75	60

Table F-4. Data Carriers in the Gaps Between Broadcasting, Amateur and CB Bands

Spectrum Other than			Data Carrier Width, kHz				
Broadcast, Amateur, CB, kHz			200	400	800	1,600	3,200
From	To	Gap	Data Carriers Available				
5,000	5,950	950	4	2	1	0	0
6,200	7,000	800	4	2	1	0	0
7,300	9,500	2,200	11	5	2	1	0
9,900	10,100	200	1	0	0	0	0
10,150	11,650	1,500	7	3	1	0	0
12,050	13,600	1,550	7	3	1	0	0
13,800	14,000	200	1	0	0	0	0
14,350	15,100	750	3	1	0	0	0
15,600	17,550	1,950	9	4	2	1	0
17,900	18,068	168	0	0	0	0	0
18,168	21,000	2,832	14	7	3	1	0
21,850	24,890	3,040	15	7	3	1	0
24,990	25,670	680	3	1	0	0	0
26,100	26,960	860	4	2	1	0	0
27,410	28,000	590	2	1	0	0	0
29,700	40,000	10,300	51	25	12	6	3
Least Gap, kHz		168					
Total Carriers			136	63	27	10	3
Total Bandwidth, kHz		28,570	27,200	25,200	21,600	16,000	9,600
Utilization, %			95	88	76	56	34

Appendix G. Example: Multiple Upstream Channels

This appendix presents an example of several upstream channels served by a single downstream channel. This is meant to illustrate one topology and one implementation of that topology.

Suppose one downstream channel is used in conjunction with four upstream channels as shown in Figure G-1. In this case, the four upstream channels are separate fibers serving four geographical communities of modems.




Figure G-1. One Downstream and Four Upstream Channels

In this topology, the CMTS transmits four Upstream Channel Descriptors (UCDs) and four MAPs. Unfortunately, each CM cannot determine to which upstream channel it is attached, because there is no way to convey the geographical information on the shared downstream channel. The CM must assume (at least at initialization) that the UCD and MAP apply to the channel to which it is attached. The CM chooses an Initial Maintenance opportunity on any of the channels and transmits a Ranging Request. The CMTS will receive the request and will redirect the CM to the appropriate upstream channel identifier by specifying the upstream channel ID in the range response. The CM MUST then use the channel ID of the range response, not the channel ID on which the range request was initiated. This is necessary only on the first ranging-response received by the CM. The CM SHOULD continue the ranging process normally and proceed to wait for station maintenance IEs. From then on, the CM will be using the MAP that is appropriate to the fiber branch to which it is connected.

A number of constraints are imposed by this topology:

- All of the upstream channels must operate at the same frequency. Since the CM is choosing a channel descriptor at random, it would be transmitting on the wrong frequency if it chose the UCD that applied to a different fiber path.
- All of the upstream channels must operate at the same symbol rate. If not, the CMTS would be unable to demodulate the Ranging Request if transmitted at the wrong symbol rate for the particular channel.
- All Initial Maintenance opportunities across all fiber branches must be aligned. When the CM randomly chooses a MAP to use, the CMTS must be prepared to receive a Ranging Request at that time.
- All Initial Maintenance opportunities must use the same burst characteristics so that the CMTS can demodulate the Ranging Request.

Note that only the initialization intervals must be aligned. Once the CM is assigned its proper channel ID, its activities need only be aligned with other users of its fiber branch. Ordinary data transmission and requests for bandwidth may occur independently across the four upstream channels.

Appendix H. The Data-Over-Cable Spanning Tree Protocol

Section 3.1.2.1 requires the use of the spanning tree protocol on CMTSs that are intended for commercial use and on bridging CMTSes. This appendix describes how the 802.1d spanning tree protocol is adapted to work for data over cable systems.

H.1 Background

A spanning tree protocol is frequently employed in a bridged network in order to deactivate redundant network connections; i.e., to reduce an arbitrary network mesh topology to an active topology that is a rooted tree that spans all of the network segments. The spanning tree algorithm and protocol should not be confused with the data-forwarding function itself; data forwarding may follow transparent learning bridge rules, or may employ any of several other mechanisms. By deactivating redundant connections, the spanning tree protocol eliminates topological loops, which would otherwise cause data packets to be forwarded forever for many kinds of forwarding devices.

A standard spanning tree protocol [IEEE 802.1d] is employed in most bridged local area networks. This protocol was intended for private LAN use and requires some modification for cable data use.

H.2 Public Spanning Tree

To use a spanning tree protocol in a public-access network such as data-over-cable, several modifications are needed to the basic IEEE 802.1d process. Primarily, the public spanning tree must be isolated from any private spanning tree networks to which it is connected. This is to protect both the public cable network and any attached private networks. Figure H-1 illustrates the general topology.




Figure H-1. Spanning Tree Topology

The task for the public spanning tree protocol, with reference to Figure H-1, is to:

- Isolate the private bridged networks from each other. If the two private networks merge spanning trees, each is subject to instabilities in the other's network. Also, the combined tree may exceed the maximum allowable bridging diameter.
- Isolate the public network from the private networks' spanning trees. The public network must not be subject to instabilities induced by customers' networks; nor should it change the spanning tree characteristics of the customers' networks.

- Disable one of the two redundant links into the cable network, so as to prevent forwarding loops. This should occur at the cable modem, rather than at an arbitrary bridge within the customer's network.

The spanning tree protocol must also serve the topology illustrated in Figure H-2:




Figure H-2. Spanning Tree Across CMTSes

In Figure H-2, in normal operation the spanning tree protocol should deactivate a link at one of the two cable modems. It should not divert traffic across the private network. Note that in some circumstances, such as deactivation of Link-X, spanning tree *will* divert traffic onto the private network (although limits on learned MAC addresses will probably throttle most transit traffic). If this diversion is undesirable, it must be prevented by means external to spanning tree; for example, by using routers.

H.3 Public Spanning Tree Protocol Details

The Data over Cable Spanning Tree algorithm and protocol is identical to that defined in [IEEE 802.1d], with the following exceptions:

- When transmitting Configuration Bridge Protocol Data Units (BPDUs), the Data over Cable Spanning Tree Multicast Address 01-E0-2F-00-00-03 is used rather than that defined in IEEE 802.1d. These BPDUs will be forwarded rather than recalculated by ordinary IEEE 802.1d bridges.
- When transmitting Configuration BPDUs, the SNAP header AA-AA-03-00-E0-2F-73-74 is used rather than the LLC 42-42-03 header employed by 802.1d. This is to further differentiate these BPDUs from those used by IEEE 802.1d bridges, in the event that some of those bridges do not correctly identify multicast MAC addresses¹¹.
- IEEE 802.1d BPDUs are ignored and silently discarded.
- Topology Change Notification (TCN) PDUs are not transmitted (or processed). TCNs are used in IEEE networks to accelerate the aging of the learning database when the network topology may have changed. Since the learning mechanism within the cable network typically differs, this message is unnecessary and may result in unnecessary flooding.
- CMTSes operating as bridges must participate in this protocol and must be assigned higher priorities (more likely to be root) than cable modems. The NSI interface on the CMTS SHOULD be assigned a port cost equivalent to a link speed of at least 100 Mbps. These two conditions, taken together, should ensure that (1) a CMTS is the root, and (2) any other CMTS will use the head-end network rather than a customer network to reach the root.
- The MAC Forwarder of the CMTS MUST forward BPDUs from upstream to downstream channels, whether or not the CMTS is serving as a router or a bridge.

Note that CMs with this protocol enabled will transmit BPDUs onto subscriber networks in order to identify other CMs on the same subscriber network. These public spanning tree BPDUs will be carried transparently over any

¹¹ It is likely that there are a number of spanning tree bridges deployed which rely solely on the LSAPs to distinguish 802.1d packets. Such devices would not operate correctly if the data-over-cable BPDUs also used LSAP=0x42.

bridged private subscriber network. Similarly, bridging CMTSes will transmit BPDUs on the NSI as well as on the RFI interface. The multicast address and SNAP header defined above are used on all links.

H.4 Spanning Tree Parameters and Defaults

Section 4.10.2 of [IEEE 802.1d] specifies a number of recommended parameter values. Those values should be used, with the exceptions listed below:

Path Cost

In [IEEE 802.1d], the following formula is used:

$$\text{Path_Cost} = 1000 / \text{Attached_LAN_speed_in_Mb/s}$$

For CMs, this formula is adapted as:

$$\text{Path_Cost} = 1000 / (\text{Upstream_symbol_rate} * \text{bits_per_symbol_for_long_data_grant})$$

That is, the modulation type (QPSK or 16QAM) for the Long Data Grant IUC is multiplied by the raw symbol rate to determine the nominal path cost. Table H-1 provides the derived values.

Table H-1. CM Path Cost

Symbol Rate	Default Path Cost	
	QPSK	16QAM
ksym/sec		
160	3125	1563
320	1563	781
640	781	391
1280	391	195
2560	195	98

For CMTSes, this formula is:

$$\text{Path_Cost} = 1000 / (\text{Downstream_symbol_rate} * \text{bits_per_symbol})$$

Bridge Priority

The Bridge Priority for CMs SHOULD default to 36864 (0x9000). This is to bias the network so that the root will tend to be at the CMTS. The CMTS SHOULD default to 32768, as per 802.1d.

Note that both of these recommendations affect only the *default* settings. These parameters, as well as others defined in 802.1d, SHOULD be manageable throughout their entire range through the Bridge MIB [RFC-1493] or other means.

Appendix I. Error Codes and Messages

These are CM and CMTS error codes and messages. These error codes are meant to emulate the standard fashion that ISDN reports error conditions regardless of the vendor producing the equipment.

The error reported are Sync loss, UCD, MAP, Ranging REQ/RSP, UCC, DHCP/TFTP failures. In some cases there is detailed error reports in other error codes are simply “it failed”.

Table I-1. Error Codes for MAC Management Messages

Error Code	Error Message
T00.0	SYNC Timing Synchronization
T01.0	Failed to acquire QAM/QPSK symbol timing. Error stats? Retry #'s?
T02.0	Failed to acquire FEC framing. Error stats? Retry #'s? # of bad frames?
T02.1	Acquired FEC framing. Failed to acquire MPEG2 Sync. Retry #'s?
T03.0	Failed to acquire MAC framing. Error stats? Retry #'s? # of bad frames?
T04.0	Failed to Receive MAC SYNC frame within time-out period.
T05.0	Loss of Sync. (Missed 5 in a row, after having SYNCed at one time)
U00.0	UCD Upstream Channel Descriptor
U01.0	No UCDs Received. Time-out
U02.0	UCD invalid or channel unusable.
U03.0	UCD valid, BUT no SYNC received. TIMED OUT.
U04.0	UCD, & SYNC valid, NO MAPS for THIS Channel.
U05.0	UCD received with invalid or out of order Configuration Change Count.
U06.0	US Channel wide parameters not set before Burst Descriptors
M00.0	MAP Upstream Bandwidth Allocation
M01.0	A transmit opportunity was missed because the MAP arrived too late.
R00.0	RNG-REQ Ranging Request
R01.0	NO Maintenance Broadcasts for Ranging opportunities Received T2 time-out.
R04.0	Received Response to Broadcast Maintenance Request, But no Unicast Maintenance opportunities received. T4 time-out.
R101.0	No Ranging Requests received from POLLED CM (CMTS generated polls)
R102.0	Retries exhausted for polled CM (report MAC address). After 16 R101.0 errors.
R103.0	Unable to Successfully Range CM (report MAC address) Retries Exhausted. Note: this is different from R102.0 in that it was able to try, i.e., got REQs but failed to Range properly.
R104.0	Failed to receive Periodic RNG-REQ from modem (SID X), timing-out SID.

R00.0	RNG-RSP Ranging Response
R02.0	No Ranging Response received, T3 time-out
R03.0	Ranging Request Retries exhausted.
R05.0	Started Unicast Maintenance Ranging no Response received. T3 time-out.
R06.0	Unicast Maintenance Ranging attempted. No Response. Retries exhausted.
R07.0	Unicast Ranging Received Abort Response. Re-initializing MAC.
I00.0	REG-REQ Registration Request
I04.0	Service not available. Reason: Other.
I04.1	Service not available. Reason: Unrecognized configuration setting.
I04.2	Service not available. Reason: Temporarily unavailable.
I04.3	Service not available. Reason: Permanent.
I101.0	Invalid MAC header.
I102.0	Invalid SID, not in use.
I103.0	Required TLVs out of order.
I104.0	Required TLVs not present.
I105.0	Downstream Frequency format invalid.
I105.1	Downstream Frequency not in use.
I105.2	Downstream Frequency invalid, not a multiple of 62500Hz.
I106.0	Upstream Channel invalid, unassigned.
I106.1	Upstream Channel Change followed with (RE-)Registration REQ.
I107.0	Upstream Channel overloaded.
I108.0	Network Access configuration has invalid parameter.
I109.0	Class of Service configuration is invalid.
I110.0	Class of Service ID unsupported.
I111.0	Class of Service ID invalid or out of range.
I112.0	Max Downstream Bit Rate configuration is invalid format.
I112.1	Max Downstream Bit Rate configuration setting is unsupported.
I113.0	Max Upstream Bit Rate configuration setting invalid format.
I113.1	Max Upstream Bit Rate configuration setting unsupported.
I114.0	Upstream Priority configuration invalid format.
I114.1	Upstream Priority configuration setting out of range.
I115.0	Guaranteed Min Upstream Channel Bit Rate configuration setting invalid format.
I115.1	Guaranteed Min Upstream Channel Bit Rate configuration setting exceeds Max Upstream Bit Rate.
I115.2	Guaranteed Min Upstream Channel Bit Rate configuration setting out of range.
I116.0	Max Upstream Channel Transmit Burst configuration setting invalid format.

I116.1	Max Upstream Channel Transmit Burst configuration setting out of range.
I117.0	Modem Capabilities configuration setting invalid format.
I117.1	Modem Capabilities configuration setting.
I00.0	REG-RSP Registration Response
I01.0	Registration RSP invalid format or not recognized
I02.0	Registration RSP not received.
I03.0	Registration RSP with bad SID.
C00.0	UCC-REQ Upstream Channel Change Request
C01.0	UCC-REQ received with invalid or out of range US channel ID.
C02.0	UCC-REQ received unable to send UCC-RSP, no TX opportunity.
C100.0	UCC-RSP Upstream Channel Change Response
C101.0	UCC-RSP not received on previous channel ID.
C102.0	UCC-RSP received with invalid channel ID.
C103.0	UCC-RSP received with invalid channel ID on new channel.
D00.0	DHCP CM Net Configuration download and Time of Day
D01.0	Discover sent no Offer received, No available DHCP Server.
D02.0	Request sent, no Response.
D03.0	Requested Info not supported.
D03.1	DHCP response doesn't contain ALL the valid fields as describe in the RF spec Appendix C.1
D04.0	Time of Day, none set or invalid data.
D04.1	Time of Day Request sent no Response received
D04.2	Time of Day Response received but invalid data/format.
D05.0	TFTP Request sent, No Response/No Server.
D06.0	TFTP Request Failed, configuration file NOT FOUND.
D07.0	TFTP Failed, OUT OF ORDER packets.
D08.0	TFTP complete, but failed Integrity Check (MIC).
B00.0	Baseline Privacy
B01.0	(For Future Study)

Appendix J. References

- [CableLabs1] Two-Way Cable Television System Characterization, Cable Television Laboratories, Inc., April 12, 1995.
- [CableLabs2] Digital Transmission Characterization of Cable Television Systems, Cable Television Laboratories, Inc., November, 1994.
- [DIX] Ethernet Protocol Version 2.0, Digital, Intel, Xerox, 1982.
- [DOCSIS3] Data-Over-Cable Interface Specifications, Cable Modem Termination System - Network Side Interface Specification, SP-CMTS-NSI-I01-960702.
- [DOCSIS4] Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, SP-CMCI-I06-010829.
- [DOCSIS5] Data-Over-Cable Service Interface Specifications 1.0, Operations Support System Interface Specification, SP-OSSI-C01-011119. See [SCTE 22-3 2002]
- [DOCSIS6] Data-Over-Cable Service Interface Specifications, Cable Modem Telephony Return Interface Specification, SP-CMTRI-I01-970804.
- [DOCSIS8] Data-Over-Cable Service Interface Specifications 1.0, Baseline Privacy Interface Specification, SP-BPI-C01-011119. See [SCTE 22-2 2002]
- [DOCSIS9] Data-Over-Cable Service Interface Specifications 1.0, Operations Support System Interface Specification Radio Frequency Interface, SP-OSSI-RFI-C01-011119. See [SCTE 22-2 2002]
- [EG 201 212] ETSI EG 201 212: "Electrical safety; Classification of interfaces for equipment to be connected to telecommunications networks". This document is also available from CENELEC as ROBT-002.
- [EN 300 429] ETSI EN 300 429 V1.2.1: "Digital Video Broadcasting (DVB); Framing structure, channel coding and modulation for cable systems", April 1998
- [EN 50081-1] CENELEC 50081-1] CENELEC 50081-1 Electromagnetic compatibility generic emission standard Part 1: Residential, commercial and light industry.
- [EN 50082-1] CENELEC 50081-1] CENELEC 50081-1 Electromagnetic compatibility generic immunity standard Part 1: Residential, commercial and light industry.
- [EN 50083-1] CENELEC EN 50083-1:1993: "Cabled distribution systems for television, sound and interactive multimedia signals, Part 1: Safety requirements".
- [EN 50083-2] CENELEC EN 50083-2: 1995 "Cabled distribution systems for television, sound and interactive multimedia signals, Part 2: Electromagnetic compatibility for equipment".
- [EN 50083-7] CENELEC EN 50083-7: "Cable networks for television signals, sound signals and interactive services, Part 7: System performance".
- [EN 50083-10] CENELEC EN 50083-10: "Cable networks for television signals, sound signals and interactive services, Part 10: System performance of return paths".
- [EN 60950] CENELEC EN 60950: "Safety of information technology equipment".
- [IEC 60169-24] Radio-frequency connectors Part 24: Radio-frequency coaxial connectors with screw coupling, typically for use in 75 ohm cable distribution systems (Type F), International Standard First edition 1991-11.

- [IEEE802] IEEE Std 802-1990, Local and Metropolitan Area Networks: Overview and Architecture.
- [IEEE802.1Q] IEEE Standard 802.1Q. Virtual Bridged Local Area Networks. December 8, 1998.
- [IS-6] EIA Interim Standard IS-6 (1983) - Recommended Cable TV Channel Identification Plan.
- [ISO8025] ISO 8025 (December 1987) - Information processing systems - Open Systems Interconnection - Specification of the Basic Encoding Rules for Abstract Syntax Notation One (ASN.1).
- [ISO8802-2] ISO/IEC 8802-2: 1994 (IEEE Std 802.2: 1994) - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 2: Logical link control
- [ISO8802-3] ISO/IEC 8802-3: 1996 (IEEE Std 802.3: 1996) - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical sublayer specifications.
- [ISO/IEC10038] ISO/IEC 10038 (ANSI/IEEE Std 802.1D): 1993, Information technology - Telecommunications and information exchange between systems - Local area networks - Media access control (MAC) bridges.
- [ISO/IEC10039] ISO/IEC 10039:1991 Information technology – Open Systems Interconnection –Local area networks –Medium Access Control (MAC) service definition.
- [ISO/IEC15802-1] ISO/IEC 10039:1991 Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 1: Medium Access Control (MAC) service definition.
- [ITU-T H.222.0] ITU-T Recommendation H.222.0 (2000) | ISO/IEC 13818-1:2000, Information technology - generic coding of moving pictures and associated audio information systems.
- [ITU-T J.83-B] Annex B to ITU-T Recommendation J.83 (4/97), Digital multi-programme systems for television sound and data services for cable distribution.
- [ITU-T X.25] ITU-T Recommendation X.25(10/96), Interface between data terminal equipment and data circuit-terminating equipment for terminals operating in the packet mode and connected to public data networks by dedicated circuit.
- [ITU-T Z.100] ITU-T Recommendation Z.100 (11/99) - CCITT Specification and description language (SDL).
- [NCTA] NCTA Recommended Practices for measurement on Cable Television Systems - National Cable Television Association, Washington DC, 2nd Edition, revised October, 1993.
- [RFC-791] Postel, J., Internet Protocol, IETF RFC-791 (MIL STD 1777), September, 1981.
- [RFC-826] Plummer, D., Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48-bit Ethernet address for transmission on Ethernet hardware, November, 1982.
- [RFC-868] Harrenstien, K., and Postel, J., Time Protocol, IETF RFC-868, May 1983.
- [RFC-1042] Postel, J., and Reynolds, J., A Standard for the Transmission of IP Datagrams over IEEE 802 Networks, IETF RFC-1042, February, 1988.
- [RFC-1058] Hedrick, C., Routing Information Protocol, IETF RFC-1058, June, 1988.
- [RFC-1123] Braden, R., Requirements for Internet Hosts -- Application and Support, IETF RFC-1123, October 1989.

- [RFC-1157] Schoffstall, M., Fedor, M., Davin, J. and Case, J., A Simple Network Management Protocol (SNMP), IETF RFC-1157, May, 1990.
- [RFC-1350] Sollins, K., The TFTP Protocol (Revision 2), IETF RFC-1350, July, 1992.
- [RFC-1493] Definitions of Managed Objects for Bridges. E. Decker, P. Langille, A. Rijssinghani, & K. McCloghrie. July 1993. (Obsoletes RFC1286)
- [RFC-1633] Braden, R., Clark, D., and Shenker, S., Integrated Services in the Internet Architecture: An Overview, IETF RFC-1633, June, 1994.
- [RFC-1812] Baker, F., Requirements for IP Version 4 Routers, IETF RFC-1812. June, 1995.
- [RFC-2104] Krawczyk, H., Bellare, M., and Canetti, R., HMAC: Keyed-Hashing for Message Authentication, IETF RFC-2104, February, 1997.
- [RFC-2131] Droms, R., Dynamic Host Configuration Protocol, IETF RFC-2131, March, 1997.
- [RFC-2132] Alexander, S., and Droms, R., DHCP Options and BOOTP Vendor Extensions, IETF RFC-2132, March, 1997.
- [RFC-2349] Malkin, G. and Harkin, A., TFTP Timeout Interval and Transfer Size Options, IETF RFC-2349, May 1998.
- [RFC-3046] Patrick, M., DHCP Relay Agent Information Option, IETF RFC-3046, January 2001.
- [SCTE 02 1997] F-connector, "F" Port (Female Indoor) Physical Dimensions [formerly IPS SP 406].
- [SCTE 22-2 2002] Data-Over-Cable Service Interface Specifications 1.0, Baseline Privacy Interface Specification, SP-BPI-C01-011119, [DOCSIS8].
- [SCTE 22-3 2002] Data-Over-Cable Service Interface Specifications 1.0, Operations Support System Interface Specification, SP-OSSI-C01-011119, [DOCSIS5] and [DOCSIS9].
- [SMS] The Spectrum Management Application (SMA) and the Common Spectrum Management Interface (csmi), Time Warner Cable, December 24, 1995.

Appendix K. Glossary

Address Resolution Protocol (ARP)

A protocol of the IETF for converting network addresses to 48-bit Ethernet addresses.

American National Standards Institute (ANSI)

A US standards body.

ANSI

See American National Standards Institute.

ARP

See Address Resolution Protocol.

Asynchronous Transfer Mode (ATM)

A protocol for the transmission of a variety of digital signals using uniform 53-byte cells.

ATM

See Asynchronous Transfer Mode.

Availability

In cable television systems, availability is the long-term ratio of the actual RF channel operation time to scheduled RF channel operation time (expressed as a percent value) and is based on a bit error rate (BER) assumption.

BPDU

See Bridge Protocol Data Unit.

Bridge Protocol Data Unit (BDU)

Spanning tree protocol messages as defined in [ISO/IEC10038].

Broadcast Addresses

A predefined destination address that denotes the set of all data network service access points.

Cable Modem (CM)

A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.

Cable Modem Termination System (CMTS)

Cable modem termination system, located at the cable television system headend or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.

Cable Modem Termination System - Network Side Interface (CMTS-NSI)

The interface, defined in [DOCSIS3], between a CMTS and the equipment on its network side.

Cable Modem to CPE Interface (CMCI)

The interface, defined in [DOCSIS4], between a CM and CPE.

Carrier Hum Modulation

The peak-to-peak magnitude of the amplitude distortion relative to the RF carrier signal level due to the fundamental and low-order harmonics of the power-supply frequency.

Carrier-to-Noise Ratio (C/N or CNR)

The square of the ratio of the root mean square (rms) of the voltage of the digitally-modulated RF carrier to the rms of the continuous random noise voltage in the defined measurement bandwidth. (If not specified explicitly, the measurement bandwidth is the symbol rate of the digital modulation; for video it is 4 MHz).

CM

See Cable Modem.

CMCI

See Cable Modem to CPE Interface.

CMTS

See Cable Modem Termination System.

CMTS-NSI

See Cable Modem Termination System - Network Side Interface.

Composite Second Order Beat (CSO)

The peak of the average level of distortion products due to second-order non-linearities in cable system equipment.

Composite Triple Beat (CTB)

The peak of the average level of distortion components due to third-order non-linearities in cable system equipment.

CPE

See Customer Premises Equipment.

Cross-Modulation

A form of television signal distortion where modulation from one or more television channels is imposed on another channel or channels.

Customer

See End User.

Customer Premises Equipment (CPE)

Equipment at the end user's premises; MAY be provided by the end user or the service provider.

Data Link Layer

Layer 2 in the Open System Interconnection (OSI) architecture; the layer that provides services to transfer data over the transmission link between open systems.

DHCP

See Dynamic Host Configuration Protocol.

Distribution Hub

A location in a cable television network which performs the functions of a Headend for customers in its immediate area, and which receives some or all of its television program material from a Master Headend in the same metropolitan or regional area.

Downstream

In cable television, the direction of transmission from the headend to the subscriber.

Dynamic Host Configuration Protocol (DHCP)

An Internet protocol used for assigning network-layer (IP) addresses.

Dynamic Range

The ratio between the greatest signal power that can be transmitted over a multichannel analog transmission system without exceeding distortion or other performance limits, and the least signal power that can be utilized without exceeding noise, error rate or other performance limits.

Electronic Industries Association (EIA)

A voluntary body of manufacturers which, among other activities, prepares and publishes standards.

End User

A human being, organization, or telecommunications system that accesses the network in order to communicate via the services provided by the network.

Errored Second

Any 1-sec interval containing at least one bit error.

Extended Subsplit

A frequency division scheme that allows bidirectional traffic on a single coaxial cable. Reverse path signals come to the headend from 5 to 42 MHz. Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit.

Fiber Node

A point of interface between a fiber trunk and the coaxial distribution.

Forward Channel

The direction of RF signal flow away from the headend toward the end user; equivalent to Downstream.

Group Delay

The difference in transmission time between the highest and lowest of several frequencies through a device, circuit or system.

Guard Time

Minimum time allocated between bursts in the upstream referenced from the symbol center of the last symbol of a burst to the symbol center of the first symbol of the following burst. The guard time should be at least the duration of five symbols plus the maximum system timing error.

Harmonic Related Carrier (HRC)

A method of spacing television channels on a cable television system in exact 6-MHz increments, with all carrier frequencies harmonically related to a common reference.

Headend

The central location on the cable network that is responsible for injecting broadcast video and other signals in the downstream direction. See also Master Headend, Distribution Hub.

Header

Protocol control information located at the beginning of a protocol data unit.

HFC

See Hybrid Fiber/Coax (HFC) System.

High Frequency (HF)

Used in this document to refer to the entire subsplit (5-30 MHz) and extended subsplit (5-42 MHz) band used in reverse channel communications over the cable television network.

Hum Modulation

Undesired modulation of the television visual carrier by the fundamental or low-order harmonics of the power supply frequency, or other low-frequency disturbances.

Hybrid Fiber/Coax (HFC) System

A broadband bidirectional 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.

ICMP

See Internet Control Message Protocol.

IEEE

See Institute of Electrical and Electronic Engineers.

IETF

See Internet Engineering Task Force.

Incremental Related Carriers (IRC)

A method of spacing NTSC television channels on a cable television system in which all channels except 5 and 6 correspond to the standard channel plan, used to reduce composite triple beat distortions.

Institute of Electrical and Electronic Engineers (IEEE)

A voluntary organization which, among other things, sponsors standards committees and is accredited by the American National Standards Institute.

International Electrotechnical Commission (IEC)

An international standards body.

International Organization for Standardization (ISO)

An international standards body, commonly known as the International Standards Organization.

Internet Control Message Protocol (ICMP)

An Internet network-layer protocol.

Internet Engineering Task Force (IETF)

A body responsible, among other things, for developing standards used in the Internet.

Impulse Noise

Noise characterized by non-overlapping transient disturbances.

Internet Protocol (IP)

An Internet network-layer protocol.

IP

See Internet Protocol.

Latency

The time, expressed in quantity of symbols, taken for a signal element to pass through a device.

Layer

A subdivision of the Open System Interconnection (OSI) architecture, constituted by subsystems of the same rank.

LLC

See Logical Link Control (LLC) procedure.

Local Area Network (LAN)

A non-public data network in which serial transmission is used for direct data communication among data stations located on the user's premises.

Logical Link Control (LLC) procedure

In a local area network (LAN) or a Metropolitan Area Network (MAN), that part of the protocol that governs the assembling of data link layer frames and their exchange between data stations, independent of how the transmission medium is shared.

MAC

See Media Access Control (MAC) procedure.

MAC Service Access Point

See Section 6.1.2.1.

Master Headend

A headend which collects television program material from various sources by satellite, microwave, fiber and other means, and distributes this material to Distribution Hubs in the same metropolitan or regional area. A Master Headend MAY also perform the functions of a Distribution Hub for customers in its own immediate area.

Media Access Control (MAC) address

The “built-in” hardware address of a device connected to a shared medium.

Media Access Control (MAC) procedure

In a subnetwork, that part of the protocol that governs access to the transmission medium independent of the physical characteristics of the medium, but taking into account the topological aspects of the subnetworks, in order to enable the exchange of data between nodes. MAC procedures include framing, error protection, and acquiring the right to use the underlying transmission medium.

Media Access Control (MAC) sublayer

The part of the data link layer that supports topology-dependent functions and uses the services of the Physical Layer to provide services to the logical link control (LLC) sublayer.

Micro-reflections

Echoes in the forward transmission path due to departures from ideal amplitude and phase characteristics.

Mini-Slot

A “mini-slot” is an integer multiple of 6.25-microsecond increments. The relationship between mini-slots, bytes and time ticks is described in Section 6.5.4.

Moving Picture Experts Group (MPEG)

A voluntary body which develops standards for digital compressed moving pictures and associated audio.

MPEG

See Moving Picture Experts Group.

MSAP

See MAC Service Access Point.

National Cable Television Association (NCTA)

A voluntary association of cable television operators which, among other things, provides guidance on measurements and objectives for cable television systems in the USA.

National Television Systems Committee (NTSC)

Committee which defined the analog color television broadcast standard used today in North America.

Network Layer

Layer 3 in the Open System Interconnection (OSI) architecture; the layer that provides services to establish a path between open systems.

Network Management

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.

Open Systems Interconnection (OSI)

A set of ISO standards for communication between different systems made by different vendors, in which the communication process is organized into seven categories in a layered sequence based on their relationship to the user. Each layer uses the layer immediately below it and provides a service to the layer above. Layers 7–4 deal with end-to-end communication between the message source and destination, and layers 3–1 with network functions.

OSI

See Open Systems Interconnection.

Packet Identifier (PID)

A unique integer value used to identify elementary streams of a program in a single- or multi-program MPEG-2 stream.

PHY

See Physical (PHY) Layer.

Physical (PHY) Layer

Layer 1 in the Open System Interconnection (OSI) architecture; the layer that provides services to transmit bits or groups of bits over a transmission link between open systems and which entails electrical, mechanical and handshaking procedures.

Physical Media Dependent (PMD) Sublayer

A sublayer of the Physical Layer which is concerned with transmitting bits or groups of bits over particular types of transmission link between open systems and which entails electrical, mechanical and handshaking procedures.

PID

See Packet Identifier.

PMD

See Physical Media Dependent (PMD) Sublayer.

Protocol

A set of rules and formats that determines the communication behavior of layer entities in the performance of the layer functions.

QAM

See Quadrature Amplitude Modulation.

QPSK

See Quadrature Phase-Shift Keying.

Quadrature Amplitude Modulation (QAM)

A method of modulating digital signals onto an RF carrier signal, using both amplitude and phase coding.

Quadrature Phase-Shift Keying (QPSK)

A method of modulating digital signals onto an RF carrier signal using four phase states to code two digital bits.

Radio Frequency (RF)

In cable television systems, this refers to electromagnetic signals in the range 5 to 1000 MHz.

Request For Comments (RFC)

A technical policy document of the IETF. See <http://ds.internic.net/ds/rfcindex.html>.

Return Loss

The parameter describing the attenuation of a guided wave signal (e.g., via a coaxial cable) returned to a source by a device or medium resulting from reflections of the signal generated by the source.

Reverse Channel

The direction of signal flow towards the headend, away from the subscriber; equivalent to Upstream.

RFC

See Request for Comments.

Service Access Point (SAP)

The point at which services are provided by one layer, or sublayer to the layer immediately above it.

Service Identifier (SID)

See Section 6.3.

SID

See Service Identifier.

Simple Network Management Protocol (SNMP)

A network management protocol of the IETF.

SMS

See Spectrum Management System.

SNAP

See Subnetwork Access Protocol.

SNMP

See Simple Network Management Protocol.

Spectrum Management System (SMS)

A system, defined in [SMS], for managing the RF cable spectrum.

Sublayer

A subdivision of a layer in the Open System Interconnection (OSI) reference model.

Subnetwork

Subnetworks are physically formed by connecting adjacent nodes with transmission links.

Subnetwork Access Protocol (SNAP)

An extension of the LLC header to accommodate the use of 802-type networks as IP networks.

Subscriber

See End User.

Subsplit

A frequency-division scheme that allows bi-directional traffic on a single cable. Reverse path signals come to the headend from 5 to 30 (up to 42 on Extended Subsplit systems) MHz. Forward path signals go from the headend from 50 or 54 MHz to the upper frequency limit of the cable network.

Subsystem

An element in a hierarchical division of an Open System that interacts directly with elements in the next higher division or the next lower division of that open system.

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.

TFTP

See Trivial File-Transfer Protocol.

Tick

A 6.25-microsecond time interval. The reference for upstream mini-slot definition and transmission times.

TLV

See Type/Length/Value.

Transit Delay

The time difference between the instant at which the first bit of a PDU crosses one designated boundary, and the instant at which the last bit of the same PDU crosses a second designated boundary.

Transmission Control Protocol (TCP)

A transport-layer Internet protocol which ensures successful end-to-end delivery of data packets without error.

Transmission Convergence Sublayer

A sublayer of the Physical Layer that provides an interface between the Data Link Layer and the PMD Sublayer.

Transmission Link

The physical unit of a subnetwork that provides the transmission connection between adjacent nodes.

Transmission Medium

The material on which information signals may be carried; e.g., optical fiber, coaxial cable, and twisted-wire pairs.

Transmission System

The interface and transmission medium through which peer physical-layer entities transfer bits.

Transport Stream

In MPEG-2, a packet-based method of multiplexing one or more digital video and audio streams having one or more independent time bases into a single stream.

Trivial File-Transfer Protocol (TFTP)

An Internet protocol for transferring files without the requirement for user names and passwords that is typically used for automatic downloads of data and software.

Type/Length/Value (TLV)

An encoding of three fields, in which the first field indicates the type of element, the second the length of the element, and the third field the value.

Upstream

The direction from the subscriber location toward the headend.

Appendix L. Revision History

SP-RFI-I01-970326 – March 26, 1997

SP-RFI-I02-971008 – October 8, 1997

SP-RFI-I03-980202 – February 2, 1998

SP-RFI-I04-980724 – July 24, 1998

SP-RFI-I05-991105 – November 5, 1999

SP-RFI-I06-010829 – August 29, 2001

SP-RFI-C01-011119 – November 19, 2001

Appendix M. Acknowledgments

John Ulm, David Unger, Wilson Sawyer, and Gerry White of Bay Networks, Inc.'s Data Over Cable Division wrote Section 3 and Sections 5 through 9, as well as Appendices A through D and G.

Steven Anderson, Tom Kolze, Victor Hou, and Clive Holborow of General Instrument Corporation wrote Section 4 and Appendix E.

CableLabs and the cable industry as a whole are grateful to these individuals and organizations for their contributions.

Appendix N. European Specification Additions

This section applies to the second technology option referred to in Section 1 (1.1 Scope). For the first option, refer to Sections 2, 4, and 5.

This appendix describes the physical-layer specifications required for what is generally called Euro-DOCSIS cable modems. This is an optional appendix and in no way affects certification of North American, DOCSIS 1.0 modems.

The numbering of the paragraphs has been made so that the suffix after N refers to the part of the specification which has changed. As a consequence some paragraphs are missing in the appendix, because no change is required.

N.2 Functional Assumptions

This section describes the characteristics of cable television plants to be assumed for the purpose of operating a data-over-cable-system. It is not a description of CMTS or CM parameters. The data-over-cable system **MUST** be interoperable with the environment described in this section.

N.2.1 Broadband access network

A coaxial-based broadband access network is assumed. This may take the form of either an all-coax or hybrid fiber/coax (HFC) network. The generic term “cable network” is used here to cover all cases.

A cable network uses a shared-medium, tree-and-branch architecture with analog transmission. The key functional characteristics assumed in this document are the following:

- Two-way transmission.
- A maximum optical/electrical spacing between the CMTS and the most distant customer terminal of 160 km.
- A maximum differential optical/electrical spacing between the CMTS and the closest and most distant modems of 160 km.

N.2.2 Equipment assumptions

N.2.2.1 Frequency plan

In the downstream direction, the cable system is assumed to have a passband with a typical lower edge between 47 and 87.5 MHz and an upper edge which is implementation-dependent but is typically in the range of 300 to 862 MHz. Within that passband, PAL/SECAM analog television signals in 7/8 MHz channels, FM-radio signals, as well as other narrowband and wideband digital signals are assumed to be present.

In the upstream direction, the cable system is assumed to have a passband with a lower edge at 5 MHz and an upper edge which is implementation-dependent but is typically in the range of 25 to 65 MHz.

N.2.2.2 Compatibility with Other Services

The CM and CMTS **MUST** coexist with the other services on the cable network. In particular,

- a) they **MUST** operate satisfactorily in the cable spectrum assigned for CMTS-CM interoperation while the balance of the cable spectrum is occupied by any combination of television and other signals; and
- b) they **MUST NOT** cause harmful interference to any other services that are assigned to the cable network in the spectrum outside of that allocated to the CM and CMTS.

N.2.2.3 Fault Isolation Impact on Other Users

As the data-over-cable system is a shared media, point-to-multipoint system, fault-isolation procedures should take into account the potential harmful impact of faults and fault-isolation procedures on numerous users of the data-over-cable and other services.

For the interpretation of harmful impact, see section N.2.2.2 above.

N.2.2.4 Cable System Terminal Devices

See Section 1.1 Scope.

N.2.3 RF Channel Assumption

The data-over-cable system, configured with at least one set of defined physical-layer parameters (e.g., modulation, forward error correction, symbol rate, etc.) from the range of configuration settings described in this specification, **MUST** be interoperable on cable networks having characteristics defined in this section in such a manner that the forward error correction provides for equivalent operation in a cable system both with and without the impaired channel characteristics described below.

N.2.3.1 Transmission downstream

The RF channel transmission characteristics of the cable network in the downstream direction assumed for the purposes of minimal operating capability are described in Table N-1. This assumes nominal analog video carrier level (peak envelope power) in a 7/8 MHz channel bandwidth. All conditions are present concurrently.

Table N-1. Assumed downstream RF channel transmission characteristics for analog TV and sound signals

Parameter	Value
Frequency range	Cable system normal downstream operating range is from 47 MHz to as high as 862 MHz. However the operating range for data communication is from 108 to 862 MHz. The use of frequencies between 108 and 136 MHz may be forbidden due to national regulation with regard to interference with aeronautical navigation frequencies.
RF channel spacing (design bandwidth)	7/8 MHz, 8 MHz channels are used for data communication
Transit delay from headend to most distant customer	≤ 0.800 ms (typically much less)
Carrier-to-noise ratio in a 8 MHz band (analog video level)	Not less than 44 dB (Note 4)
Carrier-to-interference ratio for total power (discrete and broadband ingress signals)	Not less than 52 dB within the design bandwidth
Composite triple beat distortion for analog modulated carriers	Not greater than – 57 dBc within the design bandwidth (Note 6 a)
Composite second-order distortion for analog modulated carriers	Not greater than – 57 dBc within the design bandwidth (Note 6 b)
Cross-modulation level	Under consideration
Amplitude ripple	2.5 dB in 8 MHz
Group delay ripple in the spectrum occupied by the CMTS	100 ns over frequency range 0.5 – 4.43 MHz
Micro-reflections bound for dominant echo	–10 dBc @ ≤ 0.5 μs, –15 dBc @ ≤ 1.0 μs –20 dBc @ ≤ 1.5 μs, –30 dBc @ > 1.5 μs
Carrier hum modulation	Not greater than – 46 dBc (0.5%)
Burst noise	Not longer than 25 μs at a 10 Hz average rate
Seasonal and diurnal signal level variation	8 dB
Signal level slope, 85 - 862 MHz	12 dB
Maximum analog video carrier level at the system outlet, inclusive of above signal level variation	77 dBμV (Note 6 c)
Lowest analog video carrier level at the system outlet, inclusive of above signal level variation	60 dBμV (Note 6 d)
<p>NOTE 1 – Transmission is from the headend combiner to the CM input at the customer location.</p> <p>NOTE 2 – For measurements above, the normal downstream operating frequency band (except hum), impairments are referenced to the highest-frequency PAL/SECAM carrier level.</p> <p>NOTE 3 – For hum measurements above, the normal downstream operating frequency band, a continuous-wave carrier is sent at the test frequency at the same level as the highest-frequency PAL/SECAM carrier.</p> <p>NOTE 4 – This presumes that the digital carrier is operated at analog peak carrier level. When the digital carrier is operated below the analog peak carrier level, this C/N may be less.</p> <p>NOTE 5 – Measurements methods are defined in [CENELEC 50083-7]</p> <p>NOTE 6 – For SECAM systems the following values apply</p> <p>a) Not greater than -52 dBc within the design bandwidth</p> <p>b) Not greater than -52 dBc within the design bandwidth</p> <p>c) 74 dBμV</p> <p>d) 57 dBμV</p>	

N.2.3.2 Transmission Upstream

The RF channel transmission characteristics of the cable network in the upstream direction assumed for the purposes of minimal operating capability are described in Table N-2. All conditions are at present concurrently.

Table N-2. Assumed upstream RF channel transmission characteristics

Parameter	Value
Frequency range	5 up to 65 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	≤ 0.800 ms (typically much less)
Carrier-to-noise ratio in active channel	Not less than 22 dB
Carrier-to-ingress power (the sum of discrete and broadband ingress signals) ratio in active channel	Not less than 22 dB (Note 2)
Carrier-to-interference (the sum of noise, distortion, common-path distortion and cross-modulation) ratio in active channel	Not less than 22 dB
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μs at a 1 kHz average rate for most cases (Notes 3 and 4)
Amplitude ripple	5-65 MHz: 2.5 dB in 2 MHz
Group delay ripple	5-65 MHz: 300 ns in 2MHz
Micro-reflections – Single echo	-10 dBc @ ≤ 0.5 μs -20 dBc @ ≤ 1.0 μs -30 dBc @ > 1.0 μs
Seasonal and diurnal signal level variation	Not greater than 12 dB min to max
NOTE 1 – Transmission is from the CM output at the customer location to the headend.	
NOTE 2 – Ingress avoidance or tolerance techniques MAY be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 0 dBc.	
NOTE 3 – Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier.	
NOTE 4 – Impulse noise levels more prevalent at lower frequencies (<15 MHz).	

N.2.3.2.1 Availability

Typical cable network availability is considerably greater than 99%.

N.2.4 Transmission Levels

The nominal power level of the downstream CMTS QAM signal(s) within an 8 MHz channel is targeted to be in the range -13 dBc to 0 dBc relative to the analog video carrier level and will normally not exceed the analog video carrier level (typically between -10 to -6 dBc for 64QAM, and between -6 to -4 dBc for 256QAM). The nominal power level of the upstream CM signal(s) will be as low as possible to achieve the required margin above noise and interference. Uniform power loading per unit bandwidth is commonly followed in setting upstream signal levels, with specific levels established by the cable network operator to achieve the required carrier-to-noise and carrier-to-interference ratios.

N.2.5 Frequency Inversion

There will be no frequency inversion in the transmission path in either the downstream or upstream directions, i.e., a positive change in frequency at the input to the cable network will result in a positive change in frequency at the output.

N.4 Physical Media Dependent Sublayer Specification

N.4.1 Scope

This specification defines the electrical characteristics and protocol for a Cable Modem (CM) and Cable Modem Termination System (CMTS). It is the intent of this specification to define an interoperable CM and CMTS such that any implementation of a CM can work with any CMTS. It is not the intent of this specification to imply any specific implementation.

N.4.2 Upstream

N.4.2.1 Overview

The upstream Physical Media Dependent (PMD) sublayer uses an FDMA/TDMA burst modulation format, that provides five symbol rates and two modulation formats (QPSK and 16QAM). The modulation format includes pulse shaping for spectral efficiency, is carrier-frequency agile, and has selectable output power level. The PMD sublayer format includes a variable-length modulated burst with precise timing beginning at boundaries spaced at integer multiples of 6.25 μ s apart (which is 16 symbols at the highest data rate).

Each burst supports a flexible modulation, symbol rate, preamble, randomization of the payload, and programmable FEC encoding.

All of the upstream transmission parameters associated with burst transmission outputs from the CM are configurable by the CMTS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis.

The PMD sublayer can support a near-continuous mode of transmission, wherein ramp-down of one burst MAY overlap the ramp-up of the following burst, so that the transmitted envelope is never zero. The system timing of the TDMA transmissions from the various CMs MUST provide that the centre of the last symbol of one burst and the centre of the first symbol of the preamble of an immediately following burst are separated by at least the duration of five symbols. The guard time MUST be greater than or equal to the duration of five symbols plus the maximum timing error. Timing error is contributed by both the CM and CMTS. CM timing performance is specified in sections N.4.2.7, N.4.2.8, N.4.2.10 and N.4.3.7. Maximum timing error and guard time may vary with CMTSes from different vendors.

The upstream modulator is part of the cable modem which interfaces with the cable network. The modulator contains the actual electrical-level modulation function and the digital signal-processing function; the latter provides the FEC, preamble prepend, symbol mapping, and other processing steps. This specification is written with the idea of buffering the bursts in the signal processing portion, and with the signal processing portion:

- 1) accepting the information stream a burst at a time;
- 2) processing this stream into a complete burst of symbols for the modulator; and
- 3) feeding the properly-timed bursted symbol stream to a memoryless modulator at the exact burst transmit time.

The memoryless portion of the modulator only performs pulse shaping and quadrature upconversion.

At the Demodulator, similar to the Modulator, there are two basic functional components: the demodulation function and the signal processing function. Unlike the Modulator, the Demodulator resides in the CMTS and the specification is written with the concept that there will be one demodulation function (not necessarily an actual physical

demodulator) for each carrier frequency in use. The demodulation function would receive all bursts on a given frequency.

NOTE: The unit design approach should be cognizant of the multiple-channel nature of the demodulation and signal processing to be carried out at the headend, and partition/share functionality appropriately to optimally leverage the multi-channel application. A Demodulator design supporting multiple channels in a Demodulator unit may be appropriate.

The demodulation function of the Demodulator accepts a varying-level signal centred around a commanded power level and performs symbol timing and carrier recovery and tracking, burst acquisition, and demodulation. Additionally, the demodulation function provides an estimate of burst timing relative to a reference edge, an estimate of received signal power, an estimate of signal-to-noise ratio, and may engage adaptive equalization to mitigate the effects of:

- a) echoes in the cable plant;
- b) narrowband ingress; and
- c) group delay.

The signal-processing function of the Demodulator performs the inverse processing of the signal-processing function of the Modulator. This includes accepting the demodulated burst data stream and decoding, etc., and possibly multiplexing the data from multiple channels into a single output stream. The signal-processing function also provides the edge-timing reference and gating-enable signal to the demodulators to activate the burst acquisition for each assigned burst slot. The signal-processing function may also provide an indication of successful decoding, decoding error, or fail-to-decode for each code word and the number of corrected Reed-Solomon symbols in each code word. For every upstream burst, the CMTS has a prior knowledge of the exact burst length in symbols (see N.4.2.6, N.4.2.10.1 and Appendix A.2).

N.4.2.2 Modulation Formats

The upstream modulator MUST provide both QPSK and 16QAM modulation formats.

The upstream demodulator MUST support QPSK and 16QAM modulation formats.

N.4.2.2.1 Modulation Rates

The upstream modulator MUST provide QPSK at 160, 320, 640, 1280, and 2560 ksym/sec, and 16QAM at 160, 320, 640, 1280, and 2560 ksym/sec.

This variety of modulation rates, and flexibility in setting upstream carrier frequencies, permits operators to position carriers in gaps in the pattern of narrowband ingress.

The upstream symbol rate MUST be fixed for each upstream frequency.

N.4.2.2.2 Symbol Mapping

The modulation mode (QPSK or 16QAM) is programmable. The symbols transmitted in each mode and the mapping of the input bits to the I and Q constellation MUST be as defined in Table N-3. In the table, I₁ is the MSB of the symbol map, Q₁ is the LSB for QPSK, and Q₀ is the LSB for 16QAM. Q₁ and I₀ have intermediate bit positions in 16QAM. The MSB MUST be the first bit in the serial data into the symbol mapper.

Table N-3. I/Q mapping

QAM mode	Input bit definitions
QPSK	I ₁ Q ₁
16QAM	I ₁ Q ₁ I ₀ Q ₀

The upstream QPSK symbol mapping MUST be as shown in Figure N-1.




Figure N-1. QPSK symbol mapping

The 16QAM non-inverted (Gray-coded) symbol mapping MUST be as shown in Figure N-2.




Figure N-2. 16QAM Gray-coded symbol mapping

The 16QAM differential symbol mapping MUST be as shown in Figure N-3.




Figure N-3. 16QAM differential-coded symbol mapping

If differential quadrant encoding is enabled, the currently-transmitted symbol quadrant is derived from the previously-transmitted symbol quadrant and the current input bits via Table N-4.

Table N-4. Derivation of currently-transmitted symbol quadrant

Current input bits I(1) Q(1)	Quadrant phase change	MSBs of previously transmitted symbol	MSBs for currently transmitted symbol
00	0°	11	11
00	0°	01	01
00	0°	00	00
00	0°	10	10
01	90°	11	01
01	90°	01	00
01	90°	00	10
01	90°	10	11
11	180°	11	00
11	180°	01	10
11	180°	00	11
11	180°	10	01
10	270°	11	10
10	270°	01	11
10	270°	00	01
10	270°	10	00

N.4.2.2.3 Spectral Shaping

The upstream PMD sublayer MUST support a 25% Nyquist square root raised cosine shaping.

The occupied spectrum MUST NOT exceed the channel widths shown in Table N-5.

Table N-5. Maximum channel width

Symbol rate (ksym/sec)	Channel width (kHz) a)
160	200
320	400
640	800
1280	1600
2560	3200
a) Channel width is the -30 dB bandwidth.	

N.4.2.2.4 Upstream frequency Agility and Range

The upstream PMD sublayer MUST support operation over the frequency range of 5-65 MHz edge-to-edge.

Offset frequency resolution MUST be supported having a range of ±32 kHz (increment = 1 Hz; implement within ±10 Hz).

N.4.2.2.5 Spectrum format

The upstream modulator MUST provide operation with the format $s(t) = I(t) * \cos(\omega t) - Q(t) * \sin(\omega t)$, where t denotes time and ω denotes angular frequency.

N.4.2.3 FEC Encode

N.4.2.3.1 FEC Encode modes

The upstream modulator MUST be able to provide the following selections: Reed-Solomon codes over GF(256) with $T = 1$ to 10 or no FEC coding.

The following Reed-Solomon generator polynomial MUST be supported:

$$g(x) = (x - \alpha^0)(x - \alpha^1) \dots (x - \alpha^{2T-1})$$

where the primitive element α is 0x02 hex

The following Reed-Solomon primitive polynomial MUST be supported:

$$p(x) = x^8 + x^4 + x^3 + x^2 + 1$$

The upstream modulator MUST provide codewords from a minimum size of 18 bytes (16 information bytes [k] plus two parity bytes for $T = 1$ error correction) to a maximum size of 255 bytes (k-bytes plus parity-bytes). The uncoded word size can have a minimum of one byte.

In Shortened Last Codeword mode, the CM MUST provide the last codeword of a burst shortened from the assigned length of k data bytes per codeword as described in N.4.10.1.2.

The value of T MUST be configured in response to the Upstream Channel Descriptor from the CMTS.

N.4.2.3.2 FEC Bit-to-symbol ordering

The input to the Reed-Solomon Encoder is logically a serial bit stream from the MAC layer of the CM, and the first bit of the stream MUST be mapped into the MSB of the first Reed-Solomon symbol into the encoder. The MSB of the first symbol out of the encoder MUST be mapped into the first bit of the serial bit stream fed to the Scrambler.

Note that the MAC byte-to-serial upstream convention calls for the byte LSB to be mapped into the first bit of the serial bit stream per Section 6.2.1.3.

N.4.2.4 Scrambler (Randomiser)

The upstream modulator MUST implement a scrambler (shown in Figure N-4) where the 15-bit seed value MUST be arbitrarily programmable.

At the beginning of each burst, the register is cleared and the seed value is loaded. The seed value MUST be used to calculate the scrambler bit which is combined in an XOR with the first bit of data of each burst (which is the MSB of the first symbol following the last symbol of the preamble).

The scrambler seed value MUST be configured in response to the Upstream Channel Descriptor from the CMTS.

The polynomial MUST be $x^{15} + x^{14} + 1$.




Figure N-4. Scrambler structure

N.4.2.5 Preamble prepend

The upstream PMD sublayer MUST support a variable-length preamble field that is prepended to the data after they have been randomised and Reed-Solomon encoded.

The first bit of the Preamble Pattern is the first bit into the symbol mapper (see Figure N-8), and is I1 in the first symbol of the burst (see N.4.2.4). The first bit of the Preamble Pattern is designated by the Preamble Value Offset as described in Table 6-19, Section 6.3.3.

The value of the preamble that is prepended MUST be programmable and the length MUST be 0, 2, 4, ..., or 1024 bits for QPSK and 0, 4, 8, ..., or 1024 bits for 16QAM. Thus, the maximum length of the preamble is 512 QPSK symbols or 256QAM symbols.

The preamble length and value MUST be configured in response to the Upstream Channel Descriptor message transmitted by the CMTS.

N.4.2.6 Burst profiles

The transmission characteristics are separated into three portions:

- a) Channel parameters;
- b) Burst Profile attributes; and
- c) User Unique parameters.

The Channel parameters include:

- i) the symbol rate (five rates from 160 ksym/sec to 2.56 Msym/sec in octave steps);
- ii) the centre frequency (Hz); and
- iii) the 1024-bit Preamble Superstring.

The Channel parameters are further described in Section 6.3.3, Table 6-18; these characteristics are shared by all users on a given channel. The Burst Profile attributes are listed in Table N-6, and are further described in Section 6.3.3, Table 6-19; these parameters are the shared attributes corresponding to a burst type. The User Unique Parameters may vary for each user even when using the same burst type on the same channel as another user (for example, Power Level) and are listed in Table N-7.

Table N-6. Burst Profile attributes

Burst Profile attributes	Configuration settings
Modulation	QPSK, 16QAM
Diff Enc	On/Off
Preamble Length	0-1024 bits (Note Section N.4.2.5)
Preamble Value offset	0 to 1022
FEC Error Correction (T bytes)	0 to 10 (0 implies FEC <input type="checkbox"/> off)
FEC Codeword Information Bytes (k)	Fixed: 16 to 253 (assuming FEC on) Shortened: 16 to 253 (assuming FEC on)
Scrambler Seed	15 bits
Maximum Burst Length (minislots) (Note)	0 to 255
Guard Time	5 to 255 symbols
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off
NOTE – A burst length of 0 mini-slots in the Channel Profile means that the burst length is variable on that channel for that burst type. The burst length, while not fixed, is granted explicitly by the CMTS to the CM in the MAP.	

Table N-7. User Unique Burst parameters

User Unique parameter	Configuration settings
Power Level (Note)	.8 to 55 dBmV (16QAM) .8 to 58 dBmV (QPSK) 1 –dB steps
Offset Frequency (Note)	Range <input type="checkbox"/> <input type="checkbox"/> 32 kHz; increment <input type="checkbox"/> 1 Hz; implement within <input type="checkbox"/> 10 Hz
Ranging Offset	0 to (2 ¹⁶ -1), increments of 6.25 <input type="checkbox"/> sec/64
Burst Length (mini-slots) if variable on this channel (changes burst-to-burst)	1 to 255 mini-slots
Transmit Equalizer Coefficients (Note) (advanced modems only)	Up to 64 coefficients; 4 bytes per coefficient: 2 real and 2 complex
NOTE – Values in table apply for this given channel and symbol rate.	

The CM MUST generate each burst at the appropriate time as conveyed in the mini-slot grants provided by the CMTS MAPs (see Section 6.3.4).

The CM MUST support all burst profiles commanded by the CMTS via the Burst Descriptors in the UCD (Section 6.3.3), and subsequently assigned for transmission in a MAP (Section 6.3.4).

The CM MUST implement the Offset Frequency to within 10 Hz.

Ranging Offset is the delay correction applied by the CM to the CMTS Upstream Frame Time derived at the CM, in order to synchronize the upstream transmissions in the TDMA scheme. The Ranging Offset is an advancement equal to roughly the round-trip delay of the CM from the CMTS. The CMTS MUST provide feedback correction for this offset to the CM, based on reception of one or more successfully received bursts (i.e., satisfactory result from each technique employed: error correction and/or CRC), with accuracy within 1/2 symbol and resolution of 1/64 of the frame tick increment (6.25 μ sec/64 0.09765625 μ sec 1/4 the symbol duration of the highest symbol rate 10.24 MHz⁻¹). The CMTS sends adjustments to the CM, where a negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM. CM MUST implement the correction with resolution of at most 1 symbol duration (of the symbol rate in use for a given burst), and (other than a fixed bias) with accuracy

within $\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol owing to resolution. The accuracy of CM burst timing of $\pm 0.25 \mu\text{sec}$ plus $\pm 1/2$ symbol is relative to the mini-slot boundaries derivable at the CM based on an ideal processing of the timestamp signals received from the CMTS.

The CM MUST be capable of switching burst profiles with no reconfiguration time required between bursts except for changes in the following parameters:

- 1) Output Power;
- 2) Modulation;
- 3) Symbol Rate;
- 4) Offset frequency;
- 5) Channel Frequency; and
- 6) Ranging Offset.

For Symbol Rate, Offset frequency and Ranging Offset, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol centre of one burst and the first symbol centre of the following burst. The maximum reconfiguration time of 96 symbols should compensate for the ramp down time of one burst and the ramp up time of the next burst as well as the overall transmitter delay time including the pipeline delay and optional pre-equalizer delay. For modulation type changes, the CM MUST be able to transmit consecutive bursts as long as the CMTS allocates at least 96 symbols in between the last symbol centre of one burst and the first symbol centre of the following burst. Output Power, Symbol Rate, Offset frequency, Channel Frequency and Ranging Offset MUST NOT be changed until the CM is provided sufficient time between bursts by the CMTS. Transmitted Output Power, Symbol Rate, Offset frequency, Channel Frequency and Ranging Offset MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted. The modulation MUST NOT change while more than -30 dB of any symbol's energy of the previous burst remains to be transmitted, or more than -30 dB of any symbol's energy of the next burst has been transmitted, EXCLUDING the effect of the transmit equalizer (if present in the CM). (This is to be verified with the transmit equalizer providing no filtering; delay only, if that. Note that if the CMTS has decision feedback in its equalizer, it may need to provide more than the 96 symbol gap between bursts of different modulation type which the same CM may use; this is a CMTS decision.) Negative ranging offset adjustments will cause the 96 symbol guard to be violated. The CMTS must assure that this does not happen by allowing extra guard time between bursts that is at least equal to the amount of negative ranging offset.

If Channel Frequency is to be changed, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 100 msec between the last symbol centre of one burst and the first symbol of the following burst.

The Channel Frequency of the CM MUST be settled within the phase noise and accuracy requirements of N.4.9.5 and N.4.9.6 within 100 msec from the beginning of the change.

If Output Power is to be changed by 1 dB or less, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 5 μsec between the last symbol centre of one burst and the first symbol centre of the following burst.

If Output Power is to be changed by more than 1 dB, the CM MUST be able to implement the change between bursts as long as the CMTS allocates at least 96 symbols plus 10 μsec between the last symbol centre of one burst and the first symbol centre of the following burst.

The Output Power of the CM MUST be settled to within $\pm 0.1 \text{ dB}$ of its final output power level:

- a) within 5 μsec from the beginning of a change of 1 dB or less; and
- b) within 10 μsec from the beginning of a change of greater than 1 dB.

The output transmit power MUST be maintained constant within a TDMA burst to within less than 0.1 dB (excluding the amount theoretically present due to pulse shaping, and amplitude modulation in the case of 16QAM).

N.4.2.7 Burst timing convention

Figure N-5 illustrates the nominal burst timing.




Figure N-5. Nominal burst timing

Figure N-6 indicates worst-case burst timing. In this example, burst N arrives 1.5 symbols late, and burst N + 1 arrives 1.5 symbols early, but separation of 5 symbols is maintained; 8-symbol guard band shown.




Figure N-6. Worst-case burst timing

At a symbol rate of R_s , symbols occur at a rate of one each $T_s = 1/R_s$ seconds. Ramp up and Ramp down are the spread of a symbol in the time domain beyond T_s duration owing to the symbol-shaping filter. If only one symbol was transmitted, its duration would be longer than T_s due to the shaping filter impulse response being longer than T_s . The spread of the first and last symbols of a burst transmission effectively extends the duration of the burst to longer than $N * T_s$, where N is the number of symbols in the burst.

N.4.2.8 Transmit power requirements

The upstream PMD sublayer MUST support varying the amount of transmit power. Requirements are presented for:

- 1) the range of commanded transmit power;
- 2) the step size of the power commands; and
- 3) the accuracy (actual output power compared to the commanded amount) of the response to the command.

The mechanism by which power adjustments are performed is defined in Section 9.2.4. Such adjustments MUST be within the ranges of tolerances described below.

N.4.2.8.1 Output power agility and range

The output transmit power in the design bandwidth MUST be variable over the range of 8 dBmV to 55 dBmV (16QAM), or 58 dBmV (QPSK), in 1-dB steps.

The absolute accuracy of the transmitted power MUST be ± 2 dB, and the step size accuracy ± 0.4 dB, with an allowance for hysteresis while switching in/out a step attenuator (e.g., 20 dB) in which case the accuracy requirement is relaxed to ± 1.4 dB. For example, the actual power increase resulting from a command to increase the power level by 1 dB in a CM's next transmitted burst MUST be between 0.6 and 1.4 dB.

The step resolution MUST be 1 dB or less. When a CM is commanded with finer resolution than it can implement, it MUST round to the nearest supported step size. If the commanded step is half way between two supported step sizes, the CM MUST choose the smaller step. For example, with a supported step resolution of 1 dB, a command to step ± 0.5 dB would result in no step, while a command to step ± 0.75 dB would result in a ± 1 dB step.

N.4.2.9 Fidelity requirements

N.4.2.9.1 Spurious emissions

The noise and spurious power MUST NOT exceed the levels given in Tables N-8, N-9 and N-10.

In Table N-8, In-band spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include ISI. The measurement bandwidth for In-band spurious is equal to the symbol rate (e.g., 160 kHz for 160 ksym/s).

The measurement bandwidth for the 3 (or fewer) Carrier-Related Frequency Bands (below 65 MHz) is 160 kHz, with up to three 160 kHz bands, each with no more than -47 dBc, allowed to be excluded from the "Bands within 5 to 65 MHz Transmitting Burst" specs of Table N-10.

The measurement bandwidth is also 160 kHz for the Between bursts specs of Table N-8 below 65 MHz; the Transmitting burst specs apply during the mini-slots granted to the CM (when the CM uses all or a portion of the grant), and for a mini-slot before and after the granted mini-slots. (Note that a mini-slot may be as short as 32 symbols, or 12.5 μ s at the 2.56 Msym/s rate, or as short as 200 μ s at the 160 ksym/s rate.) The Between bursts specs apply except during a used grant of mini-slots, and the mini-slot before and after the used grant.

Table N-8. Spurious emissions

Parameter	Transmitting burst	Between bursts
In-band [In-band spurious includes noise, carrier leakage, clock lines, synthesizer spurious products, and other undesired transmitter products. It does not include Inter Symbol Interference (ISI)].	-40 dBc	The greater of -72 dBc or 5 dBμV
Adjacent Band	See Table N-9	The greater of -72 dBc or 5 dBμV
3 or Fewer Carrier-Related Frequency Bands (such as second harmonic, if □ 65 MHz)	-47 dBc	The greater of -72 dBc or 5 dBμV
Bands within 5 to 65 MHz (excluding assigned channel, adjacent channels, and carrier-related channels)	See Table N-10	The greater of -72 dBc or 5 dBμV
CM Integrated Spurious Emissions Limits (all in 250 kHz, includes discrettes) 87.5 to 108 MHz	30 dBμV	5 dBμV
CM Integrated Spurious Emissions Limits (all in 4.75 MHz, includes discrettes) ^{a)} 65 to 87.5 MHz 108 to 136 MHz 136 to 862 MHz	max -40 dBc, 34 dBμV 20 dBμV 15 dBμV	34 dBμV 15 dBμV max (15 dBμV, -40 dBc ^{b)})
CM Discrete Spurious Emissions Limits ^{a)} 65 to 87.5 MHz 108 to 862 MHz	max -50 dBc, 24 dBμV 10 dBμV	24 dBμV 10 dBμV
^{a)} These spec limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST NOT be greater than 20 dBμV. ^{b)} 'dBc' is relative to the received downstream signal level. Some spurious outputs are proportional to the received signal level. ^{c)} The frequencies from 108 to 136 MHz may be forbidden due to national regulations. ^{d)} These spec limits exclude three or fewer discrete spurs. Such spurs must not be greater than 20 dBμV.		

N.4.2.9.1.1 Adjacent channel spurious emissions

Spurious emissions from a transmitted carrier may occur in an adjacent channel which could be occupied by a carrier of the same or different symbol rates. Table N-9 lists the required adjacent channel spurious emission levels for all combinations of transmitted carrier symbol rates and adjacent channel symbol rates. The measurement is performed in an adjacent channel interval that is of appropriate bandwidth and distance from the transmitted carrier based on the symbol rates of the transmitted carrier and of the carrier in the adjacent channel.

Table N-9. Adjacent channel spurious emissions

Transmitted carrier symbol rate	Specification in the interval	Measurement interval and distance from carrier edge	Adjacent channel carrier symbol rate
160 ksym/s	-45 dBc	20 to 180 kHz	160 ksym/s
	-45 dBc	40 to 360 kHz	320 ksym/s
	-45 dBc	80 to 720 kHz	640 ksym/s
	-42 dBc	160 to 1440 kHz	1280 ksym/s
	-39 dBc	320 to 2880 kHz	2560 ksym/s
All other symbol rates	-45 dBc	20 to 180 kHz	160 ksym/s
	-45 dBc	40 to 360 kHz	320 ksym/s
	-45 dBc	80 to 720 kHz	640 ksym/s
	-44 dBc	160 to 1440 kHz	1280 ksym/s
	-41 dBc	320 to 2880 kHz	2560 ksym/s

N.4.2.9.1.2 Spurious emissions in 5 to 65 MHz

Spurious emissions, other than those in an adjacent channel or carrier related emissions listed above, may occur in intervals that could be occupied by other carriers of the same or different symbol rates. To accommodate these different symbol rates and associated bandwidths, the spurious emissions are measured in an interval equal to the bandwidth corresponding to the symbol rate of the carrier that could be transmitted in that interval. This interval is independent of the current transmitted symbol rate.

Table N-10 lists the possible symbol rates that could be transmitted in an interval, the required spurious level in that interval, and the initial measurement interval at which to start measuring the spurious emissions. Measurements should start at the initial distance and be repeated at increasing distance from the carrier until the upstream band edge, 5 MHz or 65 MHz, is reached. Measurement intervals should not include carrier-related emissions.

Table N-10. Spurious emissions in 5 to 65 MHz

Possible symbol rate in this interval	Specification in the interval	Initial measurement interval and distance from carrier edge
160 ksym/s	-53 dBc	220 to 380 kHz
320 ksym/s	-50 dBc	240 to 560 kHz
640 ksym/s	-47 dBc	280 to 920 kHz
1280 ksym/s	-44 dBc	360 to 1640 kHz
2560 ksym/s	-41 dBc	520 to 3080 kHz

N.4.2.9.2 Spurious emissions during burst on/off transients

Each transmitter MUST control spurious emissions, prior to and during ramp up and during and following ramp down, before and after a burst in the TDMA scheme.

On/off spurious emissions, such as the change in voltage at the upstream transmitter output due to enabling or disabling transmission, MUST be no more than 100 mV, and such a step MUST be dissipated no faster than 2 μ s of constant slewing. This requirement applies when the CM is transmitting at ≥ 115 dB μ V or more; at backed-off transmit levels, the maximum change in voltage MUST decrease by a factor of 2 for each 6-dB decrease of power level from 115 dB μ V, down to a maximum change of 7 mV at 91 dB μ V and below. This requirement does not apply to CM power-on and power-off transients.

The slew rate limitations of 2 μ s need not be considered for DC transients of less than 7mV.

N.4.2.9.3 Symbol Error Rate (SER)

Modulator performance MUST be within 0.5 dB of theoretical SER vs. C/N (i.e., E_s/N_0), for SER as low as 10^{-6} uncoded, for QPSK and 16QAM.

The SER degradation is determined by the cluster variance caused by the transmit waveform at the output of an ideal square-root raised-cosine receive filter. It includes the effects of ISI, spurious, phase noise, and all other transmitter degradations.

Cluster SNR should be measured on a modulation analyser using a square-root raised cosine receive filter with $\alpha = 0.25$. The measured SNR MUST be better than 30 dB.

N.4.2.9.4 Filter distortion

The following requirements assume that any pre-equalization is disabled.

N.4.2.9.4.1 Amplitude

The spectral mask MUST be the ideal square root raised cosine spectrum with $\alpha = 0.25$, within the ranges given below:

$f_c - R_s/4$ Hz to $f_c + R_s/4$ Hz: -0.3 dB to 0.3 dB

$f_c - 3R_s/8$ Hz to $f_c - R_s/4$ Hz, and $f_c + R_s/4$ Hz to $f_c + 3R_s/8$ Hz: -0.5 dB to 0.3 dB

$f_c - R_s/2$ Hz and $f_c + R_s/2$ Hz: -3.5 dB to -2.5 dB

$f_c - 5R_s/8$ Hz and $f_c + 5R_s/8$ Hz: no greater than -30 dB

where f_c is the centre frequency, R_s is the symbol rate, and the spectral density is measured with a resolution bandwidth of 10 KHz or less.

N.4.2.9.4.2 Phase

$f_c - 5R_s/8$ Hz to $f_c + 5R_s/8$ Hz: Group Delay Variation MUST NOT be greater than 100 nsec.

N.4.2.9.5 Carrier phase noise

The upstream transmitter total integrated phase noise (including discrete spurious noise) MUST be less than or equal to -43 dBc summed over the spectral regions spanning 1 kHz to 1.6 MHz above and below the carrier.

N.4.2.9.6 Channel frequency accuracy

The CM MUST implement the assigned channel frequency within ± 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

N.4.2.9.7 Symbol rate accuracy

The upstream modulator MUST provide an absolute accuracy of symbol rates ± 50 parts per million over a temperature range of 0 to 40 degrees C up to five years from date of manufacture.

N.4.2.9.8 Symbol timing jitter

Peak-to-peak symbol jitter, referenced to the previous symbol zero-crossing, of the transmitted waveform, MUST be less than 0.02 of the nominal symbol duration over a 2-sec period. In other words, the difference between the maximum and the minimum symbol duration during the 2-sec period shall be less than 0.02 of the nominal symbol duration for each of the five upstream symbol rates.

The peak-to-peak cumulative phase error, referenced to the first symbol time and with any fixed symbol frequency offset factored out, MUST be less than 0.04 of the nominal symbol duration over a 0.1-sec period. In other words, the difference between the maximum and the minimum cumulative phase error during the 0.1-sec period shall be less than 0.04 of the nominal symbol duration for each of the five upstream symbol rates. Factoring out a fixed symbol frequency offset is to be done by using the computed mean symbol duration during the 0.1 sec.

N.4.2.10 Frame structure

Figure N-7 shows two examples of the frame structure: one where the packet length equals the number of information bytes in a codeword, and another where the packet length is longer than the number of information bytes in one codeword, but less than in two codewords. Example 1 illustrates the fixed codeword length mode, and Example 2 illustrates the shortened last codeword mode. These modes are defined in Section N.4.10.1.




Figure N-7. Example frame structures with flexible burst length mode

N.4.2.10.1 Codeword length

When FEC is enabled, the CM operates in either fixed-length codeword mode or with shortened-last codeword mode. The minimum number of information bytes in a codeword in either mode is 16 bytes. Shortened-last codeword mode only provides a benefit when the number of bytes in a codeword is greater than the minimum of 16 bytes.

The following descriptions apply to an allocated grant of mini-slots in both contention and non-contention regions. (Allocation of mini-slots is discussed in Section 7 of this document). The intent of the description is to define rules and conventions such that CMs request the proper number of mini-slots and the CMTS PHY knows what to expect regarding the FEC framing in both fixed codeword length and shortened last codeword modes.

N.4.2.10.1.1 Fixed codeword length

With the fixed-length codewords, after all the data are encoded, zero-fill will occur in this codeword if necessary to reach the assigned k data bytes per codeword, and zero-fill MUST continue up to the point when no additional fixed-length codewords can be inserted before the end of the last allocated mini-slot in the grant, accounting for FEC parity and guard-time symbols.

N.4.2.10.1.2 Shortened last codeword

As shown in Figure N-7, let k_1 the number of information bytes that remain after partitioning the information bytes of the burst into full-length (k burst data bytes) codewords. The value of k_1 is less than k . Given operation in a shortened last codeword mode, let k_2 the number of burst data bytes plus zero-fill bytes in the shortened last codeword. In shortened codeword mode, the CM MUST encode the data bytes of the burst (including MAC header) using the assigned codeword size (k information bytes per codeword) until:

- 1) all the data are encoded; or
- 2) a remainder of data bytes is left over which is less than k .

Shortened last codewords shall not have less than 16 information bytes, and this is to be considered when CMs make requests of mini-slots. In shortened last codeword mode, the CM MUST zero-fill data if necessary until the end of the mini-slot allocation, which in most cases will be the next mini-slot boundary, accounting for FEC parity and guard-time symbols. In many cases, only $k_{min} - k$ zero-fill bytes are necessary to fill out a mini-slot allocation with $16 \times k_{min}$ and $k_{min} \times k$. However, note the following.

More generally, the CM MUST zero-fill data until the point when no additional fixed-length codewords can be inserted before the end of the last allocated mini-slot in the grant (accounting for FEC parity and guard-time symbols), and then, if possible, a shortened last codeword of zero-fill shall be inserted to fit into the mini-slot allocation.

If, after zero-fill of additional codewords with k information bytes, there are less than 16 bytes remaining in the allocated grant of mini-slots, accounting for parity and guard-time symbols, the CM shall not create this last shortened codeword.

N.4.2.11 Signal processing requirements

The signal processing order for each burst packet type MUST be compatible with the sequence shown in Figure N-8 and MUST follow the order of steps in Figure N-9.




Figure N-8. Signal-processing sequence




Figure N-9. TDMA upstream transmission processing

N.4.2.12 Upstream demodulator input power characteristics

The maximum total input power to the upstream demodulator **MUST NOT** exceed 95 dBμV in the 5-65 MHz frequency range of operation.

The intended received power in each carrier **MUST** be within the values shown in Table N-11.

The demodulator **MUST** operate within its defined performance specifications with received bursts within ± 6 dB of the nominal commanded received power.

Table N-11. Maximum range of commanded nominal received power in each carrier

Symbol rate (ksym/s)	Maximum range (dBμV)
160	+44 to 74
320	+47 to □□
640	+50 to □̃
1280	+53 to □3
2560	+56 to □6

N.4.2.13 Upstream electrical output from the CM

The CM **MUST** output an RF modulated signal with the characteristics delineated in Table N-12.

Table N-12. Electrical output from CM

Parameter	Value
Frequency	5 to 65 MHz edge-to-edge
Level range (one channel)	□ 8 to 115 dB μ V (16QAM) □ 8 to 118 dB μ V (QPSK)
Modulation type	QPSK and 16QAM
Symbol rate (nominal)	160, 320, 640, 1280 and 2560 ksym/s
Bandwidth	200, 400, 800, 1600 and 3200 kHz
Output impedance	75 ohms
Output return loss	□ 6 dB (5-65 MHz)
Connector	[SCTE 02 1997] or [IEC 60169-24]

N.4.3 Downstream

N.4.3.1 Downstream protocol

The downstream PMD sublayer MUST conform to [EN 300 429].

Note: Any reference in this document to the transmission of television in the forward channel that is not consistent with [EN 300 429] is outside the normative scope as only [EN 300 429] is used for digital multi-program TV distribution by cable in European applications. See sec. 1 ("scope").

N.4.3.2 Interleaving

The downstream PMD sublayer MUST support the interleaver with the characteristics defined in Table N-13. This interleaver mode fully complies with [EN 300 429].

Table N-13. Interleaver characteristics

I (Number of taps)	J (Increment)	Burst protection 64QAM/256QAM	Latency 64QAM/256QAM
12	17	18 □s/14 □s	0.43 ms/0.32 ms

N.4.3.3 Downstream frequency plan

The downstream frequency plan will include all center frequencies between 112 and 858 MHz on 250 kHz increments. It is up to the operator to decide which frequencies to use to meet national and network requirements.

N.4.3.4 CMTS output electrical

The CMTS MUST output an RF modulated signal with the following characteristics defined in Table N-14.

Table N-14. CMTS output

Parameter	Value
Centre Frequency (f_c)	112 to 858 MHz \square 30 kHz
Level	Adjustable over the range 110 to 121 dB μ V
Modulation type	64QAM and 256QAM
Symbol rate (nominal)	
64QAM	6.952 Msym/s
256QAM	6.952 Msym/s
Nominal channel spacing	8 MHz
Frequency response	
64QAM	\square 15% square root raised cosine shaping
256QAM	\square 15% square root raised cosine shaping
Total discrete spurious In-band ($f_c \pm 4$ MHz)	\square -57 dBc
In-band spurious and noise ($f_c \pm 4$ MHz)	\square -46.7 dBc; where channel spurious and noise includes all discrete spurious, noise, carrier leakage, clock lines, synthesizer products, and other undesired transmitter products. Noise within ± 50 kHz of the carrier is excluded.
Adjacent channel ($f_c \pm 4.0$ MHz) to ($f_c \pm 4.75$ MHz)	\square -58 dBc in 750 kHz.
Adjacent channel ($f_c \pm 4.75$ MHz) to ($f_c \pm 12$ MHz)	\square -60.6 dBc in 7.25 MHz, excluding up to 3 spurs, each of which must be \square -60 dBc when each is measured with 10 kHz bandwidth.
Next adjacent channel ($f_c \square 12$ MHz) to ($f_c \square 20$ MHz)	Less than the greater of -63.7 dBc or 49.3 dB μ V in 8 MHz, excluding up to three discrete spurs. The total power in the spurs must be \square -60 dBc when each is measured with 10 kHz bandwidth.
Other channels (80 MHz to 1000 MHz)	\square 49.3 dB μ V in each 8 MHz channel, excluding up to three discrete spurs. The total power in the spurs must be \square -60 dBc when each is measured with 10 kHz bandwidth.
Phase noise	1 kHz-10 kHz: -33 dBc double-sided noise power 10 kHz-50 kHz: -51 dBc double-sided noise power 50 kHz-3 MHz: -51 dBc double-sided noise power
Output impedance	75 ohms
Output return loss	\square 14 dB within an output channel up to 750 MHz; \square 13 dB in an output channel above 750 MHz
Connector	[SCTE 02 1997] or [IEC 60169-24]

N.4.3.5 Downstream electrical input to CM

The CM MUST accept an RF modulated signal with the following characteristics (see Table N-15).

Table N-15. Electrical input to CM

Parameter	Value
Centre Frequency	112 to 858 MHz \pm 30 kHz
Level Range (one channel)	43 to 73 dB μ V for 64QAM 47 to 77 dB μ V for 256QAM
Modulation Type	64QAM and 256QAM
Symbol Rate (nominal)	6.952 Msym/s (64QAM) and 6.952 Msym/s (256QAM)
Bandwidth	8 MHz (15% square root raised cosine shaping for 64QAM and 15% square root raised cosine shaping for 256QAM)
Total Input Power (80-862 MHz)	\pm 90 dB μ V
Input (load) Impedance	75 ohms
Input Return Loss	\pm 6 dB (85 to 862 MHz)
Connector	[SCTE 02 1997] or [IEC 60169-24]

N.4.3.6 CM BER performance

The bit-error-rate performance of a CM MUST be as described in this section. The requirements apply to the I \pm 12, J \pm 17 mode of interleaving.

N.4.3.6.1 64QAM

N.4.3.6.1.1 64QAM CM BER Performance

Implementation loss of the CM MUST be such that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) of 25.5 dB or greater.

N.4.3.6.1.2 64QAM image rejection performance

Performance as described in N.5.6.1.1 MUST be met with analog or digital signal at 10 dBc in any portion of the RF band other than the adjacent channels.

N.4.3.6.1.3 64QAM Adjacent channel performance

Performance as described in N.5.6.1.1 MUST be met with digital signal at 0 dBc in the adjacent channels.

Performance as described in N.5.6.1.1 MUST be met with analog signal at 10 dBc in the adjacent channels.

Performance as described in N.5.6.1.1, with an additional 0.2-dB allowance, MUST be met with digital signal at 10 dBc in the adjacent channels.

N.4.3.6.2 256QAM

N.4.3.6.2.1 256QAM CM BER Performance

Implementation loss of the CM MUST be that the CM achieves a post-FEC BER less than or equal to 10^{-8} when operating at a carrier to noise ratio (E_s/N_o) as shown in Table N-16.

Table N-16. 256QAM CM BER performance

Input receive signal level	E_s/N_0
47 dB μ V to 54 dB μ V	34.5 dB
> 54 to +77 dB μ V	31.5 dB

N.4.3.6.2.2 256QAM image rejection performance

Performance as described in N.5.6.2.1 MUST be met with analog or digital signal at 10 dBc in any portion of the RF band other than the adjacent channels.

N.4.3.6.2.3 256QAM adjacent channel performance

Performance as described in N.5.6.2.1 MUST be met with analog or digital signal at 0 dBc in the adjacent channels.

Performance as described in N.5.6.2.1, with an additional 0.5-dB allowance, MUST be met with analog signal at 10 dBc in the adjacent channels.

Performance as described in N.5.6.2.1, with an additional 1.0-dB allowance, MUST be met with digital signal at 10 dBc in the adjacent channels.

N.4.3.6.2.4 Additional specifications for QAM

The following additional specifications are given for the QAM-modulation.

Parameter	Specification
I/Q Phase offset	< 1.0 °
I/Q crosstalk	<= -50 dB
I/Q Amplitude imbalance	0.05 dB max
I/Q timing skew	< 3.0 nsec

N.4.3.7 CMTS timestamp jitter

The CMTS timestamp jitter must be less than 500 nsec peak-to-peak at the output of the Downstream Transmission Convergence Sublayer. This jitter is relative to an ideal Downstream Transmission Convergence Sublayer that transfers the MPEG packet data to the Downstream Physical Media Dependent Sublayer with a perfectly continuous and smooth clock at the MPEG packet data rate. Downstream Physical Media Dependent Sublayer processing MUST NOT be considered in timestamp generation and transfer to the Downstream Physical Media Dependent Sublayer.

Thus, any two timestamps N1 and N2 ($N2 > N1$) which were transferred to the Downstream Physical Media Dependent Sublayer at times T1 and T2 respectively must satisfy the following relationship:

$$| (N2-N1)f_{\text{CMTS}} - (T2-T1) | < 500 \times 10^{-9}$$

In the equation, the value of (N2-N1) is assumed to account for the effect of rollover of the timebase counter, and T1 and T2 represent time in seconds. f_{CMTS} is the actual frequency of the CMTS master timebase and may include a fixed frequency offset from the nominal frequency of 10.24 MHz. This frequency offset is bounded by a requirement further below in this section.

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500ns allocated for jitter at the Downstream Transmission Convergence Sublayer output must be reduced by any jitter that is introduced by the Downstream Physical Media Dependent Sublayer.

The CM is expected to meet the burst timing accuracy requirements in N.4.6 when the time stamps contain this worst-case jitter.

NOTE – Jitter is the error (i.e., measured) relative to the CMTS Master Clock. (The CMTS Master Clock is the 10.24 MHz clock used for generating the timestamps.)

The CMTS 10.24 MHz Master Clock MUST have frequency stability of ± 5 ppm, drift rate $\leq 10^{-8}$ per second, and edge jitter of ≤ 10 nsec peak-to-peak (≤ 5 nsec). (The drift rate and jitter requirements on the CMTS Master Clock implies that the duration of two adjacent segments of 10240000 cycles will be within 30 nsec, due to 10 nsec jitter on each segments' duration, and 10 nsec due to frequency drift. Durations of other counter lengths also may be deduced: adjacent 1024000 segments, ≤ 21 nsec; 1024000 length segments separated by one 10240000 cycles, ≤ 30 nsec; adjacent 102400000 segments, ≤ 120 nsec. The CMTS Master Clock MUST meet such test limits in 99% or more measurements.)

N.5 Downstream transmission convergence sublayer

N.5.1 Introduction

In order to improve demodulation robustness, facilitate common receiving hardware for both video and data, and provide an opportunity for the possible future multiplexing of video and data over the PMD sublayer bitstream defined in N.4, a sublayer is interposed between the downstream PMD sublayer and the Data-Over-Cable MAC sublayer.

The downstream bitstream is defined as a continuous series of 188-byte MPEG [ITU-T H.222.0] packets. These packets consist of a 4-byte header followed by 184 bytes of payload. The header identifies the payload as belonging to the Data-Over-Cable MAC. Other values of the header may indicate other payloads. The mixture of MAC payloads and those of other services is optional and is controlled by the CMTS.

Figure N-10 illustrates the interleaving of Data-Over-Cable (DOC) MAC bytes with other digital information (digital video in the example shown).

Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = DOC	DOC MAC payload
Header = video	Digital video payload
Header = video	Digital video payload
Header = video	Digital video payload

Figure N-10. Example of interleaving MPEG packets in downstream

N.5.2 MPEG Packet format

The format of an MPEG Packet carrying EuroDOCSIS data is shown in Figure N-11. The packet consists of a 4-byte MPEG Header, a pointer_field (not present in all packets) and the EuroDOCSIS Payload.




Figure N-11. Format of an MPEG Packet

N.5.3 MPEG Header for Euro-DOCSIS Data-Over-Cable

The format of the MPEG Transport Stream Header is defined in Section 2.4 [ITU-T H.222.0]. The particular field values that distinguish Data-Over-Cable MAC streams are defined in Table N-17. Field names are from the ITU specification.

The MPEG Header consists of 4 bytes that begin the 188-byte MPEG Packet. The format of the header for use on an Euro-DOCSIS Data-Over-Cable PID is restricted to that shown in Table N-17. The header format conforms to the MPEG standard, but its use is restricted in this specification to NOT ALLOW inclusion of an adaptation_field in the MPEG packets.

Table N-17. MPEG Header format for EuroDOCSIS Data-Over-Cable packets

Field	Length (bits)	Description
sync_byte	8	0x47; MPEG Packet Sync byte.
transport_error_indicator	1	Indicates an error has occurred in the reception of the packet. This bit is reset to zero by the sender, and set to one whenever an error occurs in transmission of the packet.
payload_unit_start_indicator	1	A value of one indicates the presence of a pointer_field as the first byte of the payload (fifth byte of the packet)
transport_priority	1	Reserved; set to zero.
PID (see Note)	13	EuroDOCSIS Data-Over-Cable well-known PID (0x1FFE)
transport_scrambling_control	2	Reserved; set to '00'.
adaptation_field_control	2	'01'; use of the adaptation_field is NOT ALLOWED on the EuroDOCSIS PID.
continuity_counter	4	Cyclic counter within this PID

N.5.4 MPEG Payload for Euro-DOCSIS Data-Over-Cable

The MPEG Payload portion of the MPEG Packet will carry the Euro-DOCSIS MAC frames. The first byte of the MPEG payload will be a 'pointer_field' if the payload_unit_start_indicator (PUSI) of the MPEG Header is set.

stuff_byte

This standard defines a *stuff_byte* pattern having a value (0xFF) that is used within the Euro-DOCSIS Payload to fill any gaps between the Euro-DOCSIS MAC frames. This value is chosen as an unused value for the first byte of the Euro-DOCSIS MAC frame. The ‘FC’ byte of the MAC Header will be defined to never contain this value. (FC_TYPE = ‘11’ indicates a MAC-specific frame, and FC_PARM = ‘11111’ is not currently used and, according to this specification, is defined as an illegal value for FC_PARM.)

pointer_field

The *pointer_field* is present as the fifth byte of the MPEG packet (first byte following the MPEG header) whenever the PUSI is set to one in the MPEG header. The interpretation of the *pointer_field* is as follows:

The *pointer_field* contains the number of bytes in this packet that immediately follow the *pointer_field* that the CM decoder must skip past before looking for the beginning of an Euro-DOCSIS MAC Frame. A *pointer_field* MUST be present if it is possible to begin a Data-Over-Cable MAC Frame in the packet, and MUST point to either:

- 1) the beginning of the first MAC frame to start in the packet; or
- 2) any *stuff_byte* preceding the MAC frame.

N.5.5 Interaction with the MAC sublayer

MAC frames may begin anywhere within an MPEG packet, MAC frames may span MPEG packets, and several MAC frames may exist within an MPEG packet.

The following figures show the format of the MPEG packets that carry Euro-DOCSIS MAC frames. In all cases, the PUSI flag indicates the presence of the *pointer_field* as the first byte of the MPEG Payload.

Figure N-12 shows a MAC Frame that is positioned immediately after the *pointer_field* byte. In this case, *pointer_field* is zero, and the Euro-DOCSIS decoder will begin searching for a valid FC byte at the byte immediately following the *pointer_field*.

MPEG Header (PUSI = 1)	<i>pointer_field</i> (= 0)	MAC Frame (up to 183 bytes)	<i>stuff_byte</i> (s) (0 or more)
---------------------------	-------------------------------	--------------------------------	--------------------------------------

Figure N-12. Packet format where a MAC Frame immediately follows the *pointer_field*

Figure N-13 shows the more general case where a MAC Frame is preceded by the tail of a previous MAC Frame and a sequence of stuffing bytes. In this case, the *pointer_field* still identifies the first byte after the tail of Frame #1 (a *stuff_byte*) as the position where the decoder should begin searching for a legal MAC sublayer FC value. This format allows the multiplexing operation in the CMTS to immediately insert a MAC Frame that is available for transmission if that frame arrives after the MPEG header and *pointer_field* have been transmitted.

MPEG Header (PUSI = 1)	<i>pointer_field</i> (= M)	Tail of MAC Frame #1 (M bytes)	<i>stuff_byte</i> (s) (0 or more)	Start of MAC Frame #2
---------------------------	-------------------------------	-----------------------------------	--------------------------------------	--------------------------

Figure N-13. Packet format with MAC Frame preceded by stuffing bytes

In order to facilitate multiplexing of the MPEG packet stream carrying Euro-DOCSIS data with other MPEG-encoded data, the CMTS SHOULD NOT transmit MPEG packets with the Euro-DOCSIS PID which contain only `stuff_bytes` in the payload area. MPEG null packets SHOULD be transmitted instead. Note that there are timing relationships implicit in the Euro-DOCSIS MAC sublayer which must also be preserved by any MPEG multiplexing operation.

Figure N-14 shows that multiple MAC frames may be contained within the MPEG Packet. The MAC frames may be concatenated one after the other or be separated by an optional sequence of stuffing bytes.

MPEG Header (PUSI = 1)	pointer_field (= 0)	MAC Frame #1	MAC Frame #2	stuff_byte(s) (0 or more)	MAC Frame #3
---------------------------	------------------------	-----------------	-----------------	------------------------------	-----------------

Figure N-14. Packet format showing multiple MAC frames in a single packet

Figure N-15 shows the case where a MAC Frame spans multiple MPEG packets. In this case, the `pointer_field` of the succeeding frame points to the byte following the last byte of the tail of the first frame.

MPEG Header (PUSI = 1)	pointer_field (= 0)	stuff_byte(s) (0 or more)	Start of MAC Frame #1 (up to 183 bytes)		
MPEG Header (PUSI = 0)	Continuation of MAC Frame # 1 (184 bytes)				
MPEG Header (PUSI = 1)	pointer_field (= M)	Tail of MAC Frame #1 (M bytes)	stuff_byte(s) (0 or more)	Start of MAC Frame #2 (M bytes)	

Figure N-15. Packet format where a MAC Frame spans multiple packets

The Transmission Convergence sublayer must operate closely with the MAC sublayer in providing an accurate timestamp to be inserted into the Time Synchronization message (refer to Section 6.3.2 and Section 7.3).

N.5.6 Interaction with the Physical layer

The MPEG-2 packet stream MUST be encoded according to [EN 300 429].

N.5.7 MPEG Header synchronization and recovery

The MPEG-2 packet stream SHOULD be declared “in frame” (i.e., correct packet alignment has been achieved) when five consecutive correct sync bytes, each 188 bytes from the previous one, have been received.

The MPEG-2 packet stream SHOULD be declared “out of frame”, and a search for correct packet alignment started, when nine consecutive incorrect sync bytes are received.

The format of MAC frames is described in detail in Section 6.



This page intentionally left blank.