Data-Over-Cable Service Interface Specifications
DOCSIS 1.1

Radio Frequency Interface Specification

CM-SP-RFIv1.1-C01-050907

CLOSED
SPECIFICATION

Notice

This document is a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry in general. Neither CableLabs nor any member company is responsible for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this specification by any party. This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding its accuracy, completeness, or fitness for a particular purpose.

© Copyright 1999-2005 Cable Television Laboratories, Inc.
All rights reserved.

Comcast, Ex. 1122

Document Status Sheet

Document Control Number:	CM-SP-RFIv1.1-C01-050907			
Reference:	Radio Frequency Interface Specification			
Revision History:	I01 — First Issued Release, March 11, 1999 I02 — Second Issued Release, July 31, 1999 I03 — Third Issued Release, November 5, 1999 I04 — Fourth Issued Release, April 7, 2000 I05 — Fifth Issued Release, July 14, 2000 I06 — Sixth Issued Release, December 15, 2000 I07 — Seventh Issued Release, August 29, 2001 I08 — Eighth Issued Release, March 1, 2002 I09 — Ninth Issued Release, August 30, 2002 I10 — Tenth Issued Release, July 30, 2003 C01— Closed this specification on September 7, 2005			
Date:	September 7, 2005			
Status Code:	~~Work in Process~~	~~Draft~~	~~Issued~~	Closed
Distribution Restrictions:	~~CableLabs only~~	~~CL Reviewers~~	~~CL Vendor~~	Public

Key to Document Status Codes

Work in Process An incomplete document, designed to guide discussion and generate feedback, that may include several alternative requirements for consideration.

Draft A document in specification format considered largely complete, but lacking review by cable industry and vendors. Drafts are susceptible to substantial change during the review process.

Issued A stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and for certification testing.

Closed A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks:

DOCSIS®, eDOCSIS™, PacketCable™, CableHome™, OpenCable™, and CableLabs® are trademarks of Cable Television Laboratories, Inc.

Table of Contents

1 SCOPE AND PURPOSE 1

1.1 SCOPE 1
1.2 REQUIREMENTS 2
1.3 BACKGROUND 2
1.3.1 Service Goals 2
1.3.2 Reference Architecture 3
1.3.3 Categories of Interface Specification 5
1.3.4 Statement of Compatibility 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 8
2.2.4 Cable System Terminal Devices 8
2.3 RF CHANNEL ASSUMPTIONS 8
2.3.1 Transmission Downstream 8
2.3.2 Transmission Upstream 9
2.4 TRANSMISSION LEVELS 10
2.5 FREQUENCY INVERSION 10

3 COMMUNICATION PROTOCOLS 11

3.1 PROTOCOL STACK 11
3.1.1 CM and CMTS as Hosts 11
3.1.2 Data Forwarding Through the CM and CMTS 12
3.2 THE MAC FORWARDER 15
3.2.1 Rules for Data-Link-Layer Forwarding 15
3.3 NETWORK LAYER 16
3.3.1 Requirements for IGMP Management 16
3.4 ABOVE THE NETWORK LAYER 19
3.5 DATA LINK LAYER 19
3.5.1 LLC Sublayer 19
3.5.2 Link-Layer Security Sublayer 19
3.5.3 MAC Sublayer 19
3.6 PHYSICAL LAYER 20
3.6.1 Downstream Transmission Convergence Sublayer 20
3.6.2 PMD Sublayer 20

4 PHYSICAL MEDIA DEPENDENT SUBLAYER SPECIFICATION 21

4.1 SCOPE 21
4.2 UPSTREAM 21

4.2.1 Overview ... *21*
4.2.2 Modulation Formats ... *22*
4.2.3 FEC Encode ... *25*
4.2.4 Scrambler (Randomizer) ... *26*
4.2.5 Preamble Prepend ... *26*
4.2.6 Transmit Pre-Equalizer ... *27*
4.2.7 Burst Profiles ... *28*
4.2.8 Burst Timing Convention ... *31*
4.2.9 Transmit Power Requirements ... *32*
4.2.10 Fidelity Requirements ... *33*
4.2.11 Frame Structure ... *37*
4.2.12 Signal Processing Requirements ... *40*
4.2.13 Upstream Demodulator Input Power Characteristics ... *41*
4.2.14 Upstream Electrical Output from the CM ... *41*
4.3 DOWNSTREAM ... 41
4.3.1 Downstream Protocol ... *41*
4.3.2 Scalable Interleaving to Support Low Latency ... *42*
4.3.3 Downstream Frequency Plan ... *42*
4.3.4 CMTS Output Electrical ... *42*
4.3.5 Downstream Electrical Input to CM ... *43*
4.3.6 CM BER Performance ... *44*
4.3.7 CMTS Timestamp Jitter ... *45*

5 DOWNSTREAM TRANSMISSION CONVERGENCE SUBLAYER ... 47
5.1 INTRODUCTION ... 47
5.2 MPEG PACKET FORMAT ... 47
5.3 MPEG HEADER FOR DOCSIS DATA-OVER-CABLE ... 48
5.4 MPEG PAYLOAD FOR DOCSIS DATA-OVER-CABLE ... 48
5.5 INTERACTION WITH THE MAC SUBLAYER ... 49
5.6 INTERACTION WITH THE PHYSICAL LAYER ... 50
5.7 MPEG HEADER SYNCHRONIZATION AND RECOVERY ... 50

6 MEDIA ACCESS CONTROL SPECIFICATION ... 51
6.1 INTRODUCTION ... 51
6.1.1 Overview ... *51*
6.1.2 Definitions ... *51*
6.1.3 Future Use ... *53*
6.2 MAC FRAME FORMATS ... 53
6.2.1 Generic MAC Frame Format ... *53*
6.2.2 Packet-Based MAC Frames ... *57*
6.2.3 ATM Cell MAC Frames ... *58*
6.2.4 Reserved PDU MAC Frames ... *59*
6.2.5 MAC-Specific Headers ... *60*

6.2.6 Extended MAC Headers *65*
6.2.7 Fragmented MAC Frames *72*
6.2.8 Error-Handling *73*
6.3 MAC MANAGEMENT MESSAGES 75
6.3.1 MAC Management Message Header *75*
6.3.2 Time Synchronization (SYNC) *77*
6.3.3 Upstream Channel Descriptor (UCD) *78*
6.3.4 Upstream Bandwidth Allocation Map (MAP) *82*
6.3.5 Ranging Request (RNG-REQ) *85*
6.3.6 Ranging Response (RNG-RSP) *86*
6.3.7 Registration Request (REG-REQ) *90*
6.3.8 Registration Response (REG-RSP) *92*
6.3.9 Registration Acknowledge (REG-ACK) *95*
6.3.10 Upstream Channel Change Request (UCC-REQ) *96*
6.3.11 Upstream Channel Change Response (UCC-RSP) *97*
6.3.12 Dynamic Service Addition — Request (DSA-REQ) *98*
6.3.13 Dynamic Service Addition — Response (DSA-RSP) *100*
6.3.14 Dynamic Service Addition — Acknowledge (DSA-ACK) *103*
6.3.15 Dynamic Service Change — Request (DSC-REQ) *104*
6.3.16 Dynamic Service Change — Response (DSC-RSP) *106*
6.3.17 Dynamic Service Change — Acknowledge (DSC-ACK) *108*
6.3.18 Dynamic Service Deletion — Request (DSD-REQ) *109*
6.3.19 Dynamic Service Deletion – Response (DSD-RSP) *110*
6.3.20 Dynamic Channel Change – Request (DCC-REQ) *111*
6.3.21 Dynamic Channel Change – Response (DCC-RSP) *118*
6.3.22 Dynamic Channel Change – Acknowledge (DCC-ACK) *120*
6.3.23 Device Class Identification Request (DCI-REQ) *121*
6.3.24 Device Class Identification Response (DCI-RSP) *121*
6.3.25 Upstream Transmitter Disable (UP-DIS) MAC Management Message *122*

7 MEDIA ACCESS CONTROL PROTOCOL OPERATION 125

7.1 UPSTREAM BANDWIDTH ALLOCATION 125
7.1.1 The Allocation MAP MAC Management Message *126*
7.1.2 Information Elements *126*
7.1.3 Requests *128*
7.1.4 Information Element Feature Usage Summary *129*
7.1.5 Map Transmission and Timing *129*
7.1.6 Protocol Example *130*
7.2 SUPPORT FOR MULTIPLE CHANNELS 131
7.3 TIMING AND SYNCHRONIZATION 131
7.3.1 Global Timing Reference *132*
7.3.2 CM Channel Acquisition *132*
7.3.3 Ranging *132*

7.3.4 Timing Units and Relationships ... 133
7.4 UPSTREAM TRANSMISSION AND CONTENTION RESOLUTION ... 134
7.4.1 Contention Resolution Overview ... 134
7.4.2 Transmit Opportunities ... 136
7.4.3 CM Bandwidth Utilization ... 136
7.5 DATA LINK ENCRYPTION SUPPORT ... 137
7.5.1 MAC Messages ... 137
7.5.2 Framing ... 137

8 QUALITY OF SERVICE & FRAGMENTATION ... 139
8.1 THEORY OF OPERATION ... 139
8.1.1 Concepts ... 140
8.1.2 Object Model ... 144
8.1.3 Service Classes ... 145
8.1.4 Authorization ... 146
8.1.5 Types of Service Flows ... 147
8.1.6 Service Flows and Classifiers ... 149
8.1.7 General Operation ... 151
8.2 UPSTREAM SERVICE FLOW SCHEDULING SERVICES ... 155
8.2.1 Unsolicited Grant Service ... 155
8.2.2 Real-Time Polling Service ... 156
8.2.3 Unsolicited Grant Service with Activity Detection ... 156
8.2.4 Non-Real-Time Polling Service ... 157
8.2.5 Best Effort Service ... 157
8.2.6 Other Services ... 158
8.2.7 Parameter Applicability for Upstream Service Scheduling ... 158
8.2.8 CM Transmit Behavior ... 158
8.3 FRAGMENTATION ... 159
8.3.1 CM Fragmentation Support ... 159
8.3.2 CMTS Fragmentation Support ... 162
8.3.3 Fragmentation Example ... 163
8.4 PAYLOAD HEADER SUPPRESSION ... 167
8.4.1 Overview ... 167
8.4.2 Example Applications ... 168
8.4.3 Operation ... 168
8.4.4 Signaling ... 171
8.4.5 Payload Header Suppression Examples ... 172

9 CABLE MODEM - CMTS INTERACTION ... 175
9.1 CMTS INITIALIZATION ... 175
9.2 CABLE MODEM INITIALIZATION ... 175
9.2.1 Scanning and Synchronization to Downstream ... 177
9.2.2 Obtain Upstream Parameters ... 178

9.2.3 Message Flows During Scanning and Upstream Parameter Acquisition 180
9.2.4 Ranging and Automatic Adjustments 181
9.2.5 Device Class Identification 185
9.2.6 Establish IP Connectivity 186
9.2.7 Establish Time of Day 186
9.2.8 Transfer Operational Parameters 187
9.2.9 Registration 187
9.2.10 Baseline Privacy Initialization 192
9.2.11 Service IDs During CM Initialization 192
9.2.12 Multiple-Channel Support 193
9.3 STANDARD OPERATION 193
9.3.1 Periodic Signal Level Adjustment 193
9.3.2 Changing Upstream Burst Parameters 195
9.3.3 Changing Upstream Channels 196
9.4 DYNAMIC SERVICE 197
9.4.1 Dynamic Service Flow State Transitions 198
9.4.2 Dynamic Service Addition 208
9.4.3 Dynamic Service Change 218
9.4.4 Dynamic Service Deletion 230
9.4.5 Dynamically Changing Downstream and/or Upstream Channels 235
9.5 FAULT DETECTION AND RECOVERY 254
9.5.1 Prevention of Unauthorized Transmissions 255

10 SUPPORTING FUTURE NEW CABLE MODEM CAPABILITIES 257

10.1 DOWNLOADING CABLE MODEM OPERATING SOFTWARE 257

APPENDIX A WELL-KNOWN ADDRESSES 259

APPENDIX B PARAMETERS AND CONSTANTS 261

APPENDIX C COMMON RADIO FREQUENCY INTERFACE ENCODINGS 263

APPENDIX D CM CONFIGURATION INTERFACE SPECIFICATION 311

APPENDIX E MAC SERVICE DEFINITION 319

APPENDIX F EXAMPLE PREAMBLE SEQUENCE 329

APPENDIX G DOCSIS V1.0/V1.1 INTEROPERABILITY 331

APPENDIX H MULTIPLE UPSTREAM CHANNELS 337

APPENDIX I THE DATA-OVER-CABLE SPANNING TREE PROTOCOL 343

APPENDIX J ERROR CODES AND MESSAGES 347

APPENDIX K DOCSIS TRANSMISSION AND CONTENTION RESOLUTION 349

APPENDIX L IGMP EXAMPLE 355

APPENDIX M UNSOLICITED GRANT SERVICES 357

APPENDIX N EUROPEAN SPECIFICATION ADDITIONS 363

APPENDIX O REFERENCES 397

APPENDIX P GLOSSARY 401

APPENDIX Q ACKNOWLEDGMENTS 415

APPENDIX R REVISIONS 417

List of Figures

FIGURE 1-1. TRANSPARENT IP TRAFFIC THROUGH THE DATA-OVER-CABLE SYSTEM 2
FIGURE 1-2. DATA-OVER-CABLE REFERENCE ARCHITECTURE 4
FIGURE 3-1. PROTOCOL STACK ON THE RF INTERFACE 11
FIGURE 3-2. DATA FORWARDING THROUGH THE CM AND CMTS 12
FIGURE 3-3. EXAMPLE CONDITION FOR NETWORK LOOPS 13
FIGURE 3-4. MAC FORWARDER 15
FIGURE 4-1. QPSK SYMBOL MAPPING 23
FIGURE 4-2. 16QAM GRAY-CODED SYMBOL MAPPING 23
FIGURE 4-3. 16QAM DIFFERENTIAL-CODED SYMBOL MAPPING 24
FIGURE 4-4. SCRAMBLER STRUCTURE 26
FIGURE 4-5. TRANSMIT PRE-EQUALIZER STRUCTURE 27
FIGURE 4-6. NOMINAL BURST TIMING 31
FIGURE 4-7. WORST-CASE BURST TIMING 32
FIGURE 4-8. EXAMPLE FRAME STRUCTURES WITH FLEXIBLE BURST LENGTH MODE 37
FIGURE 4-9. SIGNAL-PROCESSING SEQUENCE 40
FIGURE 4-10. TDMA UPSTREAM TRANSMISSION PROCESSING 40
FIGURE 5-1. EXAMPLE OF INTERLEAVING MPEG PACKETS IN DOWNSTREAM 47
FIGURE 5-2. FORMAT OF AN MPEG PACKET 48
FIGURE 5-3. PACKET FORMAT WHERE A MAC FRAME IMMEDIATELY FOLLOWS THE POINTER_FIELD 49
FIGURE 5-4. PACKET FORMAT WITH MAC FRAME PRECEDED BY STUFFING BYTES 49
FIGURE 5-5. PACKET FORMAT SHOWING MULTIPLE MAC FRAMES IN A SINGLE PACKET 50
FIGURE 5-6. PACKET FORMAT WHERE A MAC FRAME SPANS MULTIPLE PACKETS 50
FIGURE 6-1. GENERIC MAC FRAME FORMAT 53
FIGURE 6-2. UPSTREAM MAC/PMD CONVERGENCE 54
FIGURE 6-3. MAC HEADER FORMAT 55
FIGURE 6-4. ETHERNET/802.3 PACKET PDU FORMAT 58
FIGURE 6-5. RESERVED PDU FORMAT 59
FIGURE 6-6. TIMING MAC HEADER 61
FIGURE 6-7. MANAGEMENT MAC HEADER 61
FIGURE 6-8. REQUEST FRAME FORMAT 62
FIGURE 6-9. FRAGMENTATION MAC HEADER FORMAT 63
FIGURE 6-10. CONCATENATION OF MULTIPLE MAC FRAMES 64
FIGURE 6-11. CONCATENATION MAC HEADER FORMAT 65
FIGURE 6-12. EXTENDED MAC FORMAT 66
FIGURE 6-13. FRAGMENTATION DETAILS 72
FIGURE 6-14. MAC HEADER AND MAC MANAGEMENT MESSAGE HEADER FIELDS 75
FIGURE 6-15. FORMAT OF PACKET PDU FOLLOWING THE TIMING HEADER 77
FIGURE 6-16. UPSTREAM CHANNEL DESCRIPTOR 78
FIGURE 6-17. TOP-LEVEL ENCODING FOR A BURST DESCRIPTOR 79
FIGURE 6-18. EXAMPLE OF UCD ENCODED TLV DATA 81
FIGURE 6-19. MAP FORMAT 82

FIGURE 6-20. MAP INFORMATION ELEMENT STRUCTURE 84
FIGURE 6-21. PACKET PDU FOLLOWING THE TIMING HEADER 85
FIGURE 6-22. RANGING RESPONSE 86
FIGURE 6-23. GENERALIZED DECISION FEEDBACK EQUALIZATION COEFFICIENTS 88
FIGURE 6-24. GENERALIZED EQUALIZER TAP LOCATION DEFINITION 89
FIGURE 6-25. EXAMPLE OF TLV DATA 89
FIGURE 6-26. REGISTRATION REQUEST 90
FIGURE 6-27. REGISTRATION RESPONSE FORMAT 92
FIGURE 6-28. REGISTRATION ACKNOWLEDGMENT 95
FIGURE 6-29. UPSTREAM CHANNEL CHANGE REQUEST 96
FIGURE 6-30. UPSTREAM CHANNEL CHANGE RESPONSE 97
FIGURE 6-31. DYNAMIC SERVICE ADDITION — REQUEST 98
FIGURE 6-32. DYNAMIC SERVICE ADDITION — RESPONSE 100
FIGURE 6-33. DYNAMIC SERVICE ADDITION — ACKNOWLEDGE 103
FIGURE 6-34. DYNAMIC SERVICE CHANGE — REQUEST 104
FIGURE 6-35. DYNAMIC SERVICE CHANGE — RESPONSE 106
FIGURE 6-36. DYNAMIC SERVICE CHANGE — ACKNOWLEDGE 108
FIGURE 6-37. DYNAMIC SERVICE DELETION — REQUEST 109
FIGURE 6-38. DYNAMIC SERVICE DELETION — RESPONSE 110
FIGURE 6-39. DYNAMIC CHANNEL CHANGE REQUEST 111
FIGURE 6-40. DYNAMIC CHANNEL CHANGE RESPONSE 118
FIGURE 6-41. DYNAMIC CHANNEL CHANGE ACKNOWLEDGE 120
FIGURE 6-42. DEVICE CLASS IDENTIFICATION REQUEST 121
FIGURE 6-43. DEVICE CLASS IDENTIFICATION RESPONSE 122
FIGURE 6-44. UP-DIS MESSAGE FORMAT 123
FIGURE 7-1. ALLOCATION MAP 125
FIGURE 7-2. PROTOCOL EXAMPLE 130
FIGURE 8-1. PROVISIONED AUTHORIZATION MODEL "ENVELOPES" 141
FIGURE 8-2. DYNAMIC AUTHORIZATION MODEL "ENVELOPES" 142
FIGURE 8-3. CLASSIFICATION WITHIN THE MAC LAYER 143
FIGURE 8-4. THEORY OF OPERATION OBJECT MODEL 145
FIGURE 8-5. REGISTRATION MESSAGE FLOW 151
FIGURE 8-6. DYNAMIC SERVICE ADDITION MESSAGE FLOW — CM INITIATED 153
FIGURE 8-7. DYNAMIC SERVICE ADDITION MESSAGE FLOW — CMTS INITIATED 153
FIGURE 8-8. CM FRAGMENTATION FLOWCHART 161
FIGURE 8-9. EXAMPLE OF FRAGMENTING A SINGLE PACKET 165
FIGURE 8-10. FRAGMENTED CONCATENATED PACKET EXAMPLE 166
FIGURE 8-11. PAYLOAD HEADER SUPPRESSION OPERATION 170
FIGURE 8-12. PAYLOAD HEADER SUPPRESSION WITH MASKING 171
FIGURE 8-13. PAYLOAD HEADER SUPPRESSION SIGNALING EXAMPLE 172
FIGURE 8-14. UPSTREAM PAYLOAD HEADER SUPPRESSION EXAMPLE 173
FIGURE 8-15. DOWNSTREAM PAYLOAD HEADER SUPPRESSION EXAMPLE 174
FIGURE 9-1. CM INITIALIZATION OVERVIEW 176

FIGURE 9-2. SDL NOTATION 177
FIGURE 9-3. OBTAINING UPSTREAM PARAMETERS 179
FIGURE 9-4. MESSAGE FLOWS DURING SCANNING AND UPSTREAM PARAMETER ACQUISITION 180
FIGURE 9-5. RANGING AND AUTOMATIC ADJUSTMENTS PROCEDURE 181
FIGURE 9-6. INITIAL RANGING - CM 182
FIGURE 9-7. INITIAL RANGING - CM (CONTINUED) 183
FIGURE 9-8. INITIAL RANGING - CMTS 184
FIGURE 9-9. DEVICE CLASS IDENTIFICATION 185
FIGURE 9-10. ESTABLISHING IP CONNECTIVITY 186
FIGURE 9-11. ESTABLISHING TIME OF DAY 186
FIGURE 9-12. REGISTRATION — CM 188
FIGURE 9-13. WAIT FOR REGISTRATION RESPONSE — CM 189
FIGURE 9-14. REGISTRATION — CMTS 191
FIGURE 9-15. REGISTRATION ACKNOWLEDGMENT— CMTS 192
FIGURE 9-16. PERIODIC RANGING - CMTS 194
FIGURE 9-17. PERIODIC RANGING - CM VIEW 195
FIGURE 9-18. CHANGING UPSTREAM CHANNELS: CMTS VIEW 196
FIGURE 9-19. CHANGING UPSTREAM CHANNELS: CM VIEW 197
FIGURE 9-20. DYNAMIC SERVICE FLOW OVERVIEW 198
FIGURE 9-21. DYNAMIC SERVICE FLOW STATE TRANSITION DIAGRAM 201
FIGURE 9-22. DSA—LOCALLY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 202
FIGURE 9-23. DSA—REMOTELY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 203
FIGURE 9-24. DSC—LOCALLY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 204
FIGURE 9-25. DSC—REMOTELY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 205
FIGURE 9-26. DSD—LOCALLY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 206
FIGURE 9-27. DYNAMIC DELETION (DSD) - REMOTELY INITIATED TRANSACTION STATE TRANSITION DIAGRAM 207
FIGURE 9-28. DYNAMIC SERVICE ADDITION INITIATED FROM CM 208
FIGURE 9-29. DYNAMIC SERVICE ADDITION INITIATED FROM CMTS 209
FIGURE 9-30. DSA—LOCALLY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 210
FIGURE 9-31. DSA—LOCALLY INITIATED TRANSACTION DSA-RSP PENDING STATE FLOW DIAGRAM 211
FIGURE 9-32. DSA—LOCALLY INITIATED TRANSACTION HOLDING STATE FLOW DIAGRAM 212
FIGURE 9-33. DSA—LOCALLY INITIATED TRANSACTION RETRIES EXHAUSTED STATE FLOW DIAGRAM ..213
FIGURE 9-34. DSA—LOCALLY INITIATED TRANSACTION DELETING SERVICE FLOW STATE FLOW DIAGRAM 214
FIGURE 9-35. DSA—REMOTELY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 215
FIGURE 9-36. DSA—REMOTELY INITIATED TRANSACTION DSA-ACK PENDING STATE FLOW DIAGRAM 216
FIGURE 9-37. DSA—REMOTELY INITIATED TRANSACTION HOLDING DOWN STATE FLOW DIAGRAM 217
FIGURE 9-38. DSA—REMOTELY INITIATED TRANSACTION DELETING SERVICE STATE FLOW DIAGRAM ..218
FIGURE 9-39. CM-INITIATED DSC 220
FIGURE 9-40. CMTS-INITIATED DSC 221
FIGURE 9-41. DSC—LOCALLY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 222
FIGURE 9-42. DSC—LOCALLY INITIATED TRANSACTION DSC-RSP PENDING STATE FLOW DIAGRAM 223

FIGURE 9-43. DSC—LOCALLY INITIATED TRANSACTION HOLDING DOWN STATE FLOW DIAGRAM 224
FIGURE 9-44. DSC—LOCALLY INITIATED TRANSACTION RETRIES EXHAUSTED STATE FLOW DIAGRAM . 225
FIGURE 9-45. DSC—LOCALLY INITIATED TRANSACTION DELETING SERVICE FLOW STATE FLOW DIAGRAM 226
FIGURE 9-46. DSC—REMOTELY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 227
FIGURE 9-47. DSC—REMOTELY INITIATED TRANSACTION DSC-ACK PENDING STATE FLOW DIAGRAM 228
FIGURE 9-48. DSC—REMOTELY INITIATED TRANSACTION HOLDING DOWN STATE FLOW DIAGRAM 229
FIGURE 9-49. DSC—REMOTELY INITIATED TRANSACTION DELETING SERVICE FLOW STATE FLOW DIAGRAM 229
FIGURE 9-50. DYNAMIC SERVICE DELETION INITIATED FROM CM 230
FIGURE 9-51. DYNAMIC SERVICE DELETION INITIATED FROM CMTS 230
FIGURE 9-52. DSD—LOCALLY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 231
FIGURE 9-53. DSD—LOCALLY INITIATED TRANSACTION DSD-RSP PENDING STATE FLOW DIAGRAM ... 232
FIGURE 9-54. DSD—LOCALLY INITIATED TRANSACTION HOLDING DOWN STATE FLOW DIAGRAM 233
FIGURE 9-55. DSD—REMOTELY INITIATED TRANSACTION BEGIN STATE FLOW DIAGRAM 234
FIGURE 9-56. DSD—REMOTELY INITIATED TRANSACTION HOLDING DOWN STATE FLOW DIAGRAM 235
FIGURE 9-57. DCC EXAMPLE OPERATIONAL FLOW 243
FIGURE 9-58. DYNAMICALLY CHANGING CHANNELS: CMTS VIEW PART 1 245
FIGURE 9-59. DYNAMICALLY CHANGING CHANNELS: CMTS VIEW PART 2 246
FIGURE 9-60. DYNAMICALLY CHANGING CHANNELS: CMTS VIEW PART 3 247
FIGURE 9-61. DYNAMICALLY CHANGING CHANNELS: CMTS VIEW PART 4 248
FIGURE 9-62. DYNAMICALLY CHANGING CHANNELS: CM VIEW PART 1 249
FIGURE 9-63. DYNAMICALLY CHANGING CHANNELS: CM VIEW PART 2 250
FIGURE 9-64. DYNAMICALLY CHANGING CHANNELS: CM VIEW PART 3 251
FIGURE 9-65. DYNAMICALLY CHANGING CHANNELS: CM VIEW PART 4 252
FIGURE H-1. SINGLE DOWNSTREAM AND SINGLE UPSTREAM CHANNELS PER CM 338
FIGURE H-2. MULTIPLE DOWNSTREAM AND MULTIPLE UPSTREAM CHANNELS PER CM 340

List of Tables

TABLE 2-1. ASSUMED DOWNSTREAM RF CHANNEL TRANSMISSION CHARACTERISTICS (SEE NOTE 1) 9
TABLE 2-2. ASSUMED UPSTREAM RF CHANNEL TRANSMISSION CHARACTERISTICS (SEE NOTE 1) 10
TABLE 4-1. I/Q MAPPING 23
TABLE 4-2. DERIVATION OF CURRENTLY TRANSMITTED SYMBOL QUADRANT 24
TABLE 4-3. MAXIMUM CHANNEL WIDTH 25
TABLE 4-4. BURST PROFILE ATTRIBUTES 28
TABLE 4-5. USER UNIQUE BURST PARAMETERS 29
TABLE 4-6. SPURIOUS EMISSIONS 34
TABLE 4-7. ADJACENT MCHANNEL SPURIOUS EMISSIONS RELATIVE TO THE TRANSMITTED BURST POWER LEVEL 35
TABLE 4-8. SPURIOUS EMISSIONS IN 5 TO 42 MHZ RELATIVE TO THE TRANSMITTED BURST POWER LEVEL 35
TABLE 4-9. FILTER AMPLITUDE DISTORTION 36
TABLE 4-10. MAXIMUM RANGE OF COMMANDED NOMINAL RECEIVE POWER IN EACH CARRIER 41
TABLE 4-11. ELECTRICAL OUTPUT FROM CM 41
TABLE 4-12. INTERLEAVER CHARACTERISTICS 42
TABLE 4-13. CMTS OUTPUT 43
TABLE 4-14. ELECTRICAL INPUT TO CM 44
TABLE 5-1. MPEG HEADER FORMAT FOR DOCSIS DATA-OVER-CABLE PACKETS 48
TABLE 6-1. GENERIC MAC HEADER FORMAT 56
TABLE 6-2. FC FIELD FORMAT 56
TABLE 6-3. PACKET PDU FORMAT 58
TABLE 6-4. RESERVED PDU FORMAT 60
TABLE 6-5. MAC-SPECIFIC HEADERS AND FRAMES 60
TABLE 6-6. TIMING MAC HEADER FORMAT 61
TABLE 6-7. MAC MANAGEMENT FORMAT 62
TABLE 6-8. REQUEST FRAME (REQ) FORMAT 63
TABLE 6-9. FRAGMENTATION MAC FRAME (FRAG) FORMAT 64
TABLE 6-10. CONCATENATED MAC FRAME FORMAT 65
TABLE 6-11. EXAMPLE EXTENDED HEADER FORMAT 67
TABLE 6-12. EH ELEMENT FORMAT 67
TABLE 6-13. EXTENDED HEADER TYPES 68
TABLE 6-14. FRAGMENTATION EXTENDED HEADER FORMAT 69
TABLE 6-15. PAYLOAD HEADER SUPPRESSION EHDR SUB-ELEMENT FORMAT 70
TABLE 6-16. UNSOLICITED GRANT SYNCHRONIZATION EHDR SUB-ELEMENT FORMAT 71
TABLE 6-17. MAC MANAGEMENT MESSAGE TYPES 76
TABLE 6-18. CHANNEL TLV PARAMETERS 79
TABLE 6-19. UPSTREAM PHYSICAL-LAYER BURST ATTRIBUTES 80
TABLE 6-20. ALLOCATION MAP INFORMATION ELEMENTS (IE) 84
TABLE 6-21. RANGING RESPONSE MESSAGE ENCODINGS 87
TABLE 7-1. IE FEATURE COMPATIBILITY SUMMARY 129
TABLE 7-2. EXAMPLE RELATING MINI-SLOTS TO TIME TICKS 134

TABLE 7-3. TRANSMIT OPPORTUNITY 136
TABLE 8-1. TFTP FILE CONTENTS 151
TABLE 8-2. REGISTRATION REQUEST CONTENTS 152
TABLE 8-3. REGISTRATION RESPONSE CONTENTS 152
TABLE 8-4. PARAMETER APPLICABILITY FOR UPSTREAM SERVICE SCHEDULING 158
TABLE 8-5. PAYLOAD HEADER SUPPRESSION DEFINITIONS 167
TABLE 9-1. USE CHANNEL DIRECTLY: CMTS SUPPLIES ALL TLV HINTS 253
TABLE 9-2. STATION MAINTENANCE: CMTS SUPPLIES NO TLV HINTS 254
TABLE 9-3. RECOVERY PROCESS ON LOSS OF SPECIFIC MAC MESSAGES 255
TABLE 8-1. MAC MESSAGES WITH CHANNEL IDS 341

Radio Frequency Interface Specification

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 inter-operable. One technology option is based on the downstream multi-programme television distribution that is deployed in North America using 6 MHz channelling, and supports upstream transmission in the region 5-42 MHz. The other technology option is based on the corresponding European multi-programme 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 Section 2, Section 4, Section 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 Section 2, Section 4, and Section 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], [EN 50082-1], [EN 50083-2], [EN 50083-7], and [EN 50083-10] are relevant to the European Community.

The option of Sections 2, 4 and 5 together with Appendix G and Appendix C.1.1.1 is required to be backwards compatible with an earlier version of that technology [SCTE1], whereas the option of Appendix N was not included in [SCTE1] and therefore is not required to be backwards compatible with [SCTE1].

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 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 packet-based communications systems on cable television systems that are capable of supporting a wide variety of services. Services under consideration by cable operators include packet telephony service, video conferencing service, T1/frame relay equivalent service, and many others. To this end, CableLabs' member companies have decided to prepare a series of interface specifications that will 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.

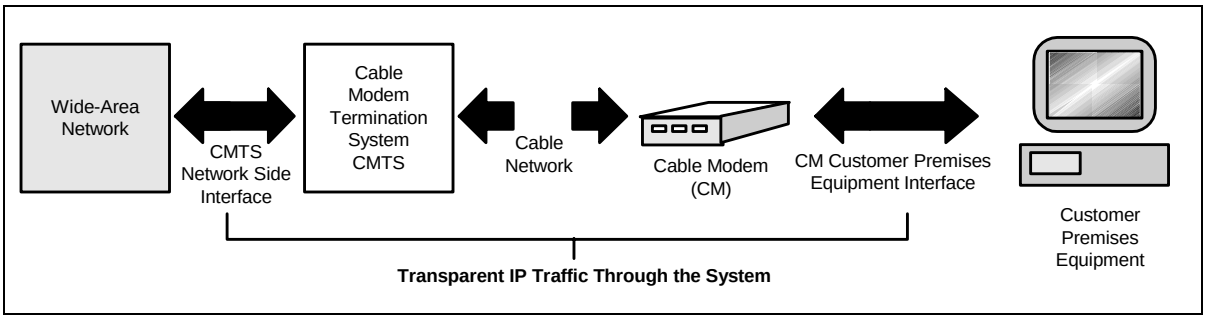


Figure 1-1. Transparent IP Traffic Through the Data-Over-Cable System

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 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

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

Note: This architecture illustrates the North American frequency plans only and is not normative for European applications. Refer to Section 1.1 for applicability.

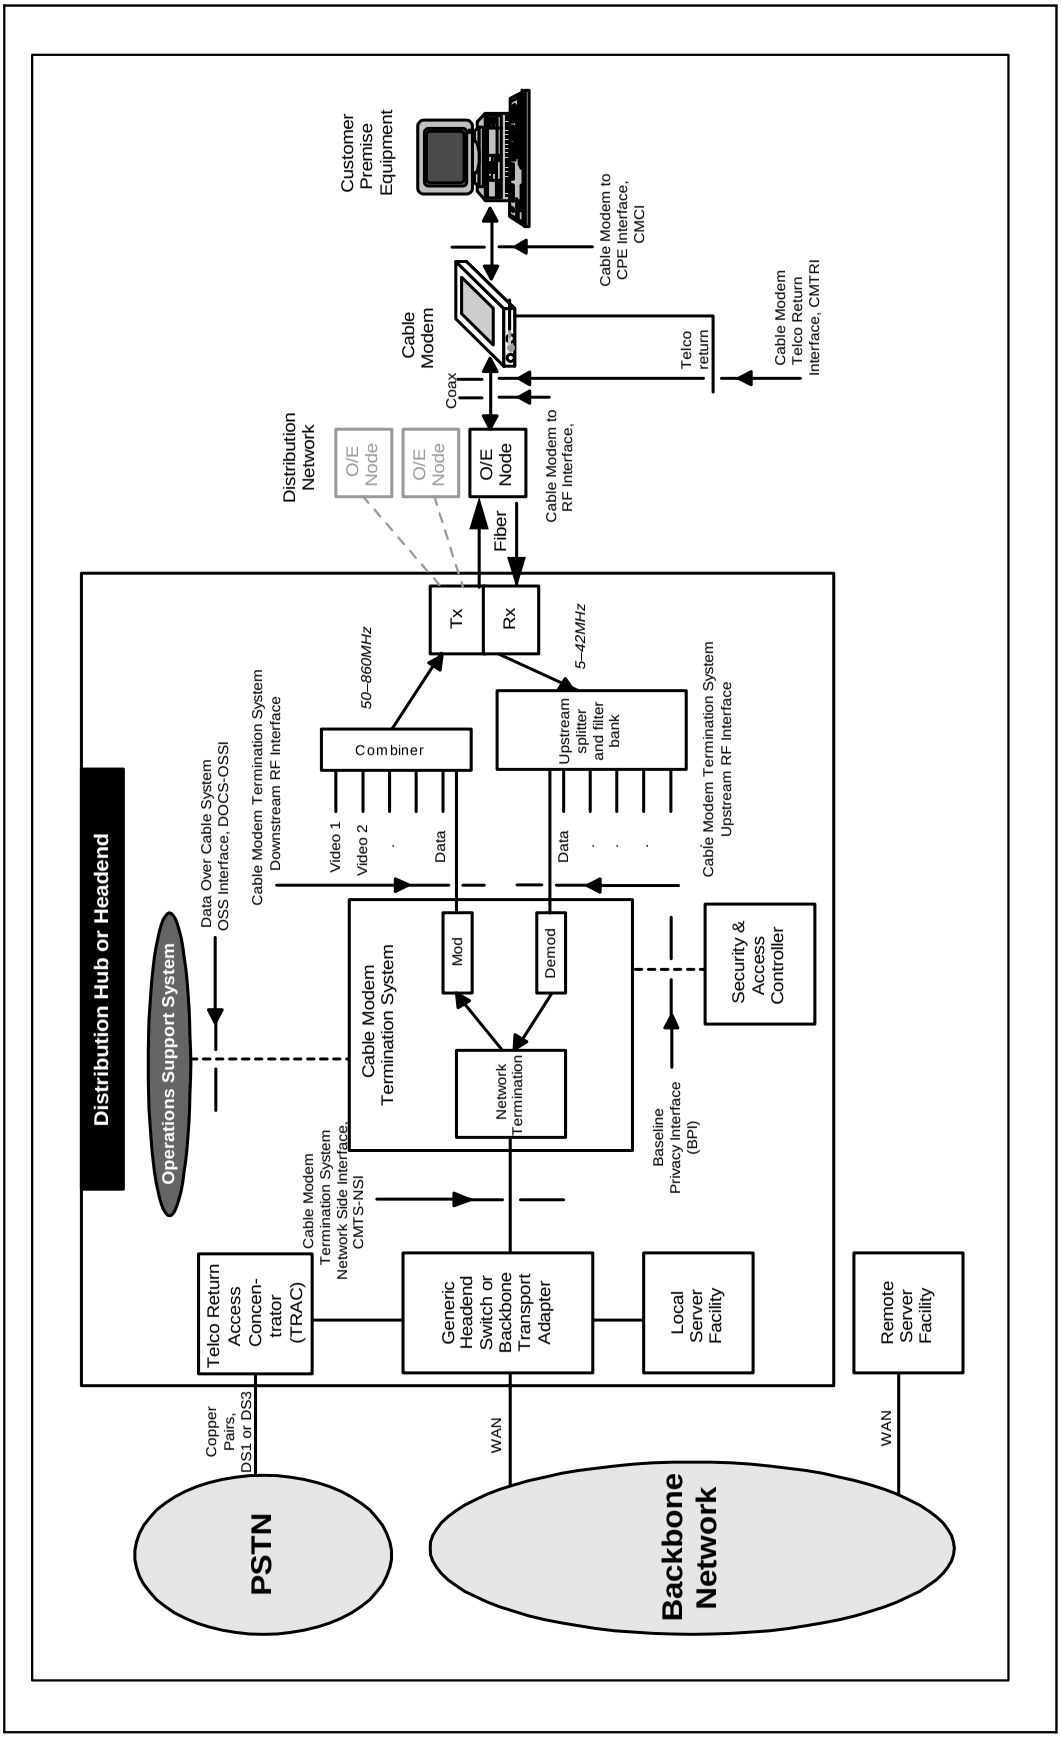


Figure 1-2. Data-Over-Cable Reference Architecture

1.3.3 Categories of Interface Specification

The basic reference architecture of Figure 1-2 involves five categories of interface.

Data Interfaces - These are the CMCI [DOCSIS4] and CMTS-NSI [DOCSIS3], corresponding respectively to the cable-modem-to-customer-premises-equipment (CPE) interface (for example, between the customer's computer and the cable modem), and the cable modem termination system network-side interface between the cable modem termination system and the data network.

Operations Support Systems Interfaces - These are network element management layer interfaces between the network elements and the high-level OSSs (operations support systems) which support the basic business processes, and are documented in [DOCSIS5].

Telephone Return Interface - CMTRI - This is the interface between the cable modem and a telephone return path, for use in cases where the return path is not provided or not available via the cable network, and is documented in [DOCSIS6].

RF Interfaces - The RF interfaces defined in this document are the following:

- Between the cable modem and the cable network
- Between the CMTS and the cable network, in the downstream direction (traffic toward the customer)
- Between the CMTS and the cable network, in the upstream direction (traffic from the customer)

Security Interfaces -

- Baseline data-over-cable security is defined in [DOCSIS8].

1.3.3.1 Data-Over-Cable Service Interface Documents

A list of the documents in the Data-Over-Cable Service Interface Specifications family is provided below. For updates, please refer to URL http://www.cablemodem.com.

Designation	Title
SP-CMCI	Cable Modem to Customer Premises Equipment Interface Specification
SP-CMTS-NSI	Cable Modem Termination System Network Side Interface Specification
SP-CMTRI	Cable Modem Telco Return Interface Specification
SP-OSSI	Operations Support System Interface Specification
SP-RFI	Radio Frequency Interface Specification
SP-BPI+	Baseline Privacy Plus Interface Specification

Key to Designations:

SP	Specification
TP	Test Plan — a document of test procedures to validate specification conformance, interoperability or performance
TR	Technical Report (provides a context for understanding and applying the specification or initial ideas about possible future features)

1.3.4 Statement of Compatibility

This document specifies an interface, commonly referred to as DOCSIS® 1.1, which is an extension of the interface specified in [SCTE1], commonly referred to as DOCSIS 1.0. These extensions are entirely backwards and forwards compatible with the previous specification. DOCSIS 1.1-compliant CMs MUST interoperate seamlessly with DOCSIS 1.0 CMTSs. DOCSIS 1.1-compliant CMTSs MUST seamlessly support DOCSIS 1.0 CMs.

Refer to Appendix G for further interoperability information.

2 Functional Assumptions

This section describes the characteristics of cable television plant 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 within the environment described in this section.

This section applies to the first technology option referred to in Section 1.1 (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 analogue 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 CM 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 between 50 and 54 MHz and an upper edge that is implementation-dependent but is typically in the range of 300 to 864 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-S542], 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 (5-40 or 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,

a) They MUST be interoperable 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 spectrum outside of that allocated to the CMTS.

The latter is understood as

- No measurable degradation (highest level of compatibility),
- No degradation below the perceptible level of impairments for all services (standard or medium level of compatibility), or
- No degradation below the minimal standards accepted by the industry (for example, FCC for analog video services) or other service provider (minimal level of compatibility).

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 2.2.2 above.

2.2.4 Cable System Terminal Devices

The CM MUST meet and SHOULD exceed all applicable regulations for Cable System Termination Devices and Cable Ready Consumer Equipment as defined in FCC Part 15 [FCC15] and Part 76 [FCC76]. None of these specific requirements may be used to relax any of the specifications contained elsewhere within this document.

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, 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.

2.3.1 Transmission Downstream

The RF channel transmission characteristics of the cable network in the downstream direction are described in Table 2-1. These numbers assume total average power of a digital signal in a 6-MHz channel bandwidth for carrier levels unless indicated otherwise. For impairment levels, the numbers in Table 2-1 assume average power in a bandwidth in which the impairment levels are measured in a standard manner for cable TV system. For analog signal levels, the numbers in Table 2-1 assume peak envelope power in a 6-MHz channel bandwidth. All conditions are present concurrently. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this specification.

Table 2-1. Assumed Downstream RF Channel Transmission Characteristics (see note 1)

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 >= 88 MHz.
RF channel spacing (design bandwidth)	6 MHz
Transit delay from headend to most distant customer	<= 0.800 msec (typically much less)
Carrier-to-noise ratio in a 6-MHz band	Not less than 35 dB (see note 2,3)
Carrier-to-Composite triple beat distortion ratio	Not less than 41 dB (see note 2,3)
Carrier-to-Composite second order distortion ratio	Not less than 41 dB (see note 2,3)
Carrier-to-Cross-modulation ratio	Not less than 41 dB (see note 2,3)
Carrier-to-any other discrete interference (ingress)	Not less than 41 dB (see note 2,3)
Amplitude ripple	3 dB within the design bandwidth (see note 2)
Group delay ripple in the spectrum occupied by the CMTS	75 ns within the design bandwidth (see note 2)
Micro-reflections bound for dominant echo	-20 dBc @ <= 1.5 µsec, -30 dBc @ > 1.5 µsec -10 dBc @ <= 0.5 µsec, -15 dBc @ <= 1.0 µsec (see note 2)
Carrier hum modulation	Not greater than -26 dBc (5%) (see note 2)
Burst noise	Not longer than 25 µsec at a 10 Hz average rate (see note 2)
Maximum analog video carrier level at the CM input	17 dBmV
Maximum number of analog carriers	121

Notes to Table 2-1:

1. Transmission is from the headend combiner to the CM input at the customer location.
2. Measurement methods defined in [NCTA] or [CableLabs2].
3. Measured relative to a QAM signal that is equal to the nominal video level in the plant.

2.3.2 Transmission Upstream

The RF channel transmission characteristics of the cable network in the upstream direction are described in Table 2-2. All conditions are present concurrently. No combination of the following parameters will exceed any stated interface limit defined elsewhere in this specification.

Table 2-2. Assumed Upstream RF Channel Transmission Characteristics (see note 1)

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-interference plus ingress (the sum of noise, distortion, common-path distortion and cross-modulation and the sum of discrete and broadband ingress signals, impulse noise excluded) ratio	Not less than 25 dB (Note 2)
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 and 4)
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 reverse gain (loss) variation	Not greater than 14 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 10 dBc. The ratios are guaranteed only within the digital carrier channels.
3. Amplitude and frequency characteristics sufficiently strong to partially or wholly mask the data carrier.
4. 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 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 CMTSs 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.

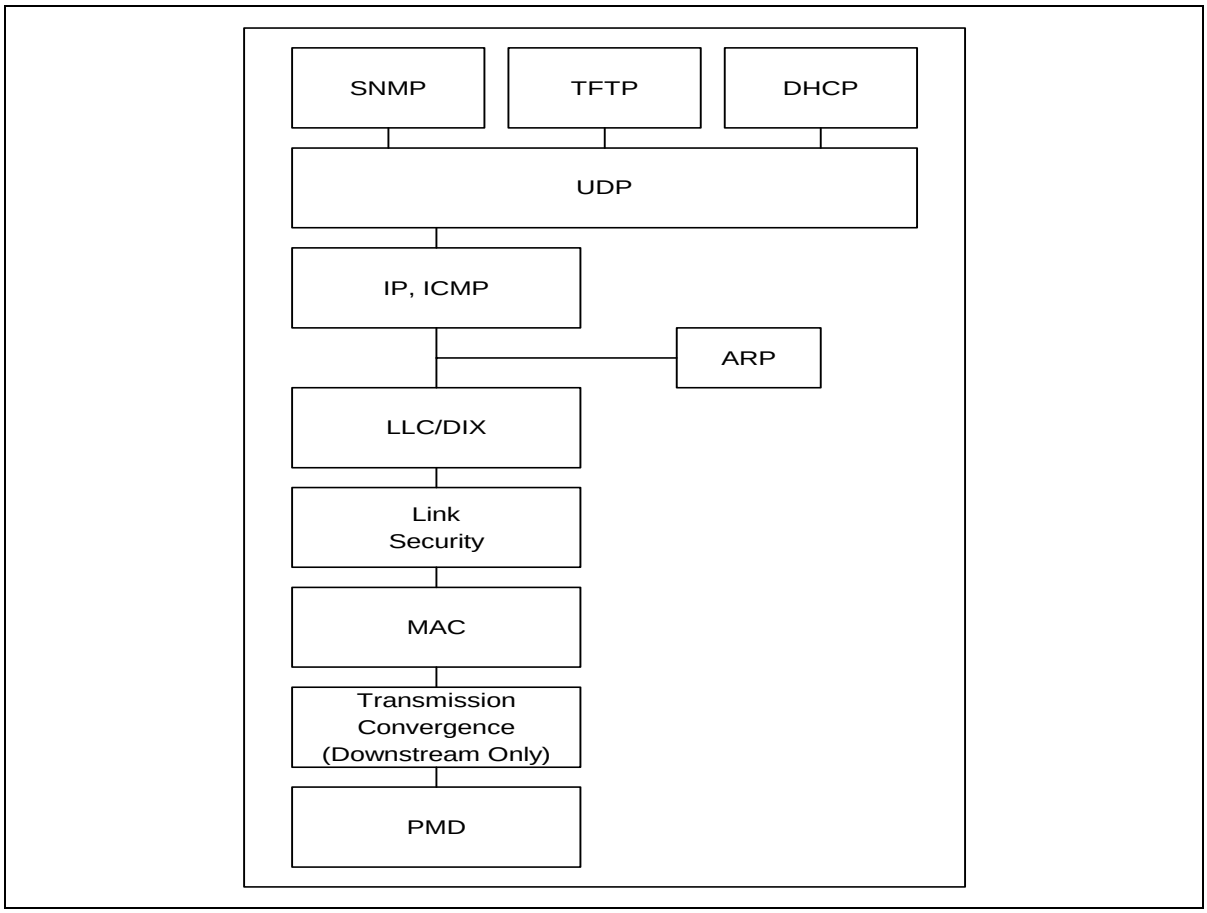


Figure 3-1. Protocol Stack on the RF Interface

The CM and CMTS MUST function as IP hosts. As such, the CM and CMTS MUST support IP and ARP over DIX link-layer framing (see [DIX]). The CMTS MUST NOT transmit frames that are smaller than the DIX 64 byte minimum on a downstream channel.[1] However, the CM MAY transmit frames that are smaller than the DIX 64 byte minimum on an upstream channel.

The CM and CMTS MAY also support IP and ARP over SNAP framing [RFC-1042].

The CM and CMTS also MUST function as LLC hosts. As such, the CM and CMTS MUST respond appropriately to TEST and XID requests per [ISO8802-2].

3.1.2 Data Forwarding Through the CM and CMTS

3.1.2.1 General

Data forwarding through the CMTS MAY be transparent bridging,[2] or MAY employ network-layer forwarding (routing, IP switching) as shown in Figure 3-2.

Data forwarding through the CM is link-layer transparent bridging, as shown in Figure 3-2. Forwarding rules are similar to [ISO/IEC10038] with the modifications described in Section 3.1.2.2 and Section 3.1.2.3. This allows the support of multiple network layers.

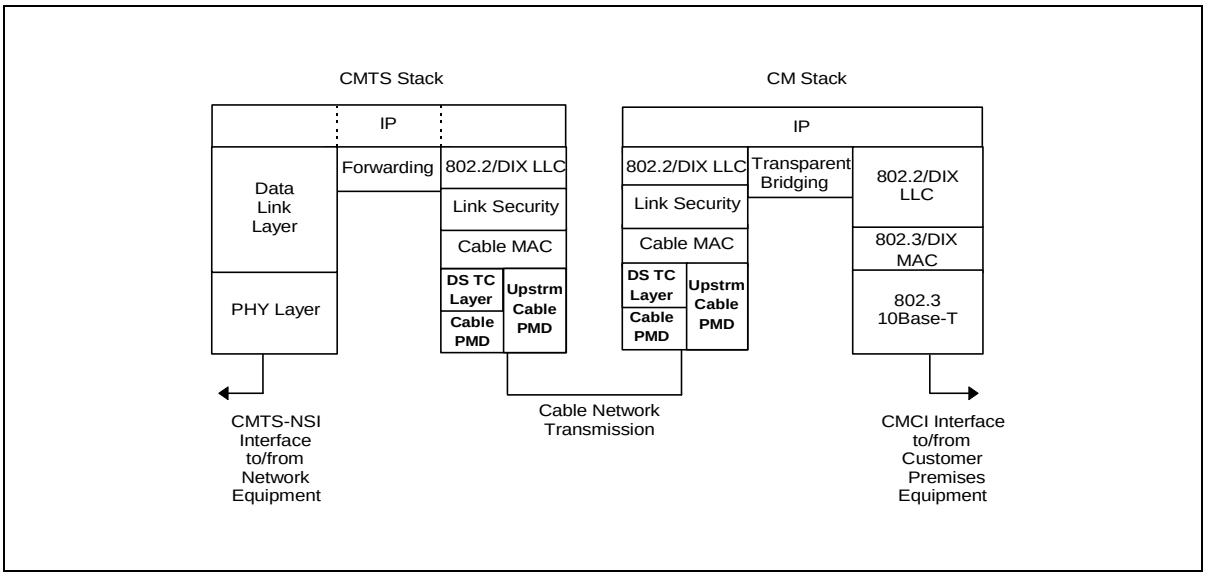


Figure 3-2. Data Forwarding Through the CM and CMTS

Forwarding of IP traffic MUST be supported. Other network layer protocols MAY be supported. The ability to restrict the network layer to a single protocol such as IP MUST be supported.

The 802.1d spanning tree protocol of [ISO/IEC10038] with the modifications described in Appendix I MAY be supported by CMs intended for residential use. CMs intended for commercial use MUST support this version of spanning tree. CMs and CMTSs MUST include the ability to filter (and disregard) 802.1d BPDUs.

---

1. Except as a result of Payload Header Suppression. Refer to Section 8.4.

2. With the exception that for packet PDUs less than 64 bytes to be forwarded from the upstream RFI, a CMTS MUST pad out the packet PDU and recompute the CRC.

This specification assumes that CMs intended for residential use will not be connected in a configuration which would create network loops such as that shown in Figure 3-3.

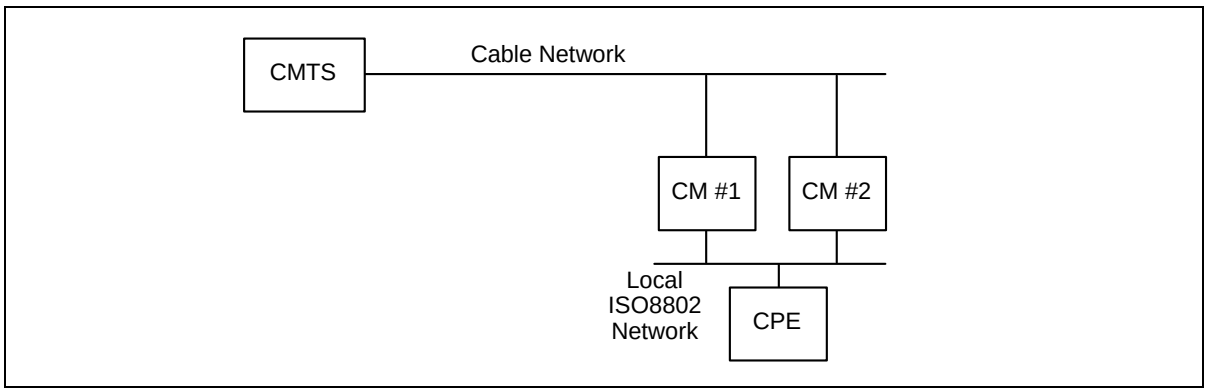


Figure 3-3. Example Condition for Network Loops

Although provisions exist in this specification for frames to be passed from a higher-layer entity to be forwarded by the cable modem, these frames MUST be treated identically to frames arriving at the CPE port. In particular, all of the forwarding rules defined in Section 3.1.2.3 MUST apply to these frames.

3.1.2.2 CMTS Forwarding Rules

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

- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.
- Link-layer frames, on a given Service Flow (refer to Section 6.1.2.3), MUST be delivered in the order they are received.

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

If network-layer forwarding is used, then 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.
- 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 MAC addresses (a device-dependent value). Once the CM acquires its maximum number of CPE MAC addresses, then newly discovered CPE MAC addresses MUST NOT replace previously acquired addresses. The CM must support acquisition of at least one CPE MAC 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 precedence over learned addresses.
- CPE addresses MUST NOT be aged out.
- In order to allow modification of user MAC addresses or movement of the CM, addresses are not retained in non-volatile storage. On a CM reset (*e.g.*, power cycle), all provisioned and learned addresses MUST be discarded.

3.1.2.3.2 Forwarding

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

- Link-layer frames MUST NOT be duplicated.
- Stale frames (those that cannot be delivered in a timely fashion) MUST be discarded.
- Link-layer frames MUST be delivered in the order that they are received on a given Service Flow (refer to Section 6.1.2.5). In the upstream direction, the CM may perform one or more frame/packet processing functions on frames received from the CMCI prior to classifying them to a Service Flow. In the downstream direction, the CM may perform one or more frame/packet processing functions on frames received from the HFC prior to transmitting them on the CMCI. Example processing functions include: DOCSIS protocol filtering as specified in [DOCSIS5] section 4.3, a policy-based filtering service as described in Section 8.1.6.1 and Appendix E, and priority-based queuing to support 802.1P/Q services.

Cable-Network-to-CMCI forwarding MUST follow the following specific rules:

- Frames addressed to unknown destinations MUST NOT be forwarded from the cable port to the CPE ports.
- Broadcast frames MUST be forwarded to the CPE ports, unless they are from source addresses which are provisioned or learned as supported CPE devices, in which case they MUST NOT be forwarded.
- The forwarding of multicast is controlled by administratively set parameters for the policy-filter service and by a specific multicast tracking algorithm (refer to Section 3.3.1). Multicast frames MUST NOT be forwarded unless both mechanisms are in a permissive state.

CMCI-to-Cable-Network forwarding MUST follow the following specific rules:

- Frames addressed to unknown destinations MUST be forwarded from all CPE ports 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.

- Other (non-supported) CPE source addresses MUST be learned from all CPE ports and this information used to filter local traffic as in a traditional learning bridge.
- Frames addressed to destination addresses that are learned from all CPE ports MUST be filtered as local traffic.

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.

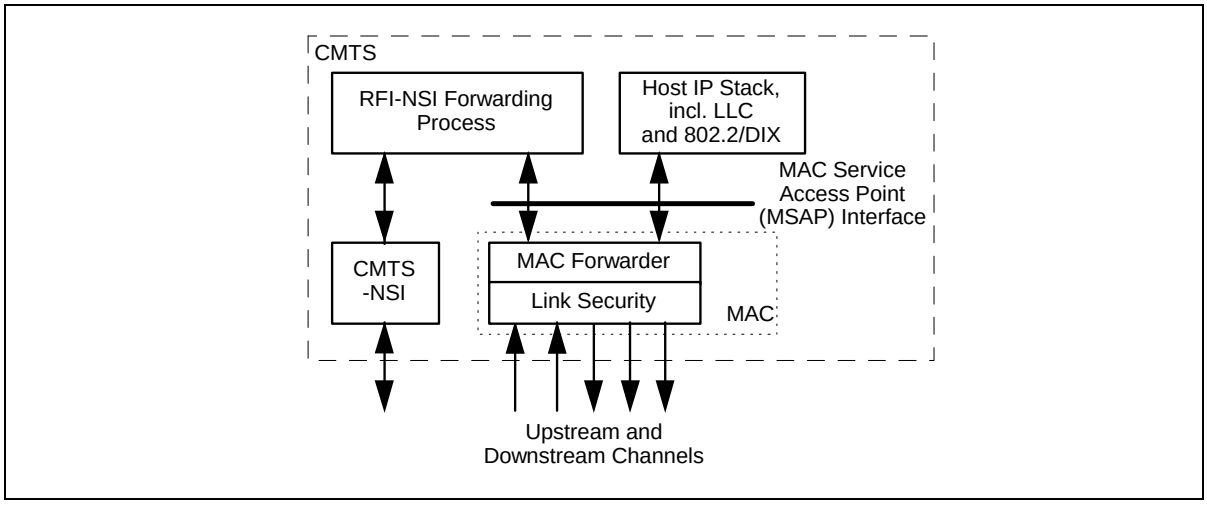


Figure 3-4. MAC Forwarder

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 Rules for Data-Link-Layer Forwarding

If the MAC Forwarder is implemented using only data-link-layer semantics, then 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, then the frame MUST be forwarded to that channel.[3]

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, then the frame MUST be delivered to the MSAP interface.

If the destination address is broadcast, multicast[4], or unknown, the frame MUST be delivered to both the MSAP and to all downstream channels. (With the exception of the 3.3.1.2 multicast forwarding rules.)

Delivery rules are similar to those for transparent bridging:

- 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.
- Frames, on a given Service Flow (refer to Section 6.1.2.3), MUST be delivered in the order they are received.

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.

3.3.1 Requirements for IGMP Management

There are two basic modes of IGMP capability that are applicable to a DOCSIS 1.1 device (CMTS and CM). The first mode is a *passive* operation in which the device selectively forwards IGMP based upon the known state of multicast session activity on the subscriber side (an example of this is described in Appendix L). In *passive* mode, the device derives its IGMP timers based on the rules specified in Section 3.3.1.1 of the RFI. The second mode is an *active* operation in which the device terminates and initiates IGMP based upon the known state of multicast session activity on the subscriber side. One example of the latter, active, mode is commonly referred to as an IGMP-Proxy implementation side (as described in [ID-IGMP]). A more complete example of an active IGMP device is that of a Multicast Router.

Active and Passive IGMP devices MUST support IGMPv2 [RFC-2236]

---

[3] 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.

[4] All multicasts, including 802.1d/ISO 10038 Spanning Tree Bridge BPDU's, MUST be forwarded.

3.3.1.1 IGMP Timer Requirements

The following IGMP timer requirements apply only when the device (CMTS / CM) is operating in passive IGMP mode:

- The device MUST NOT require any specific configuration for the associated multicast timer values and MUST be capable of adhering to the timers specified in this section.
- The device MAY provide configuration control that overrides the default values of these timers.
- The device MUST derive the Membership Query Interval by looking at the inter-arrival times of the Membership Query messages. Formally: If $n < 2$, MQI = 125 else MQI = MAX (125, MQ_n- MQ_{n-1}), where MQI is the Membership Query Interval in seconds, n is the number of Membership Queries seen, and MQ_n is the epoch time at which the nth Membership Query was seen to the nearest second.
- The Query Response Interval is carried in the Membership Query packet. The Query Response Interval MUST be assumed to be 10 seconds if not otherwise set (or set to 0) in the Membership Query packet.

3.3.1.2 CMTS Rules

- If link-layer forwarding of multicast packets is used, the CMTS MUST forward all Membership Queries on all downstream channels using the appropriate 802.3 multicast group (*e.g.*, 01:00:5E:xx:xx:xx where xx:xx:xx are the low order 23 bits of the multicast address expressed in hex notation. Refer to [IMA].).
- The CMTS MUST forward the first copy of Solicited and Unsolicited Membership Reports for any given group received on its upstream RF interface to all of its downstream RF interfaces. However, if membership is managed on a per downstream RF interface basis, Membership Reports and IGMP v2 Leave messages MAY be forwarded only on the downstream interface to which the reporting CPE's CM is connected.
- The CMTS SHOULD suppress the transmission of additional Membership Reports (for any given group) downstream for at least the Query Response Interval. If the CMTS uses data-link-layer forwarding, it MUST also forward the Membership Report out all appropriate Network Side Interfaces.
- The CMTS SHOULD suppress the downstream transmission of traffic to any IP multicast group that does not have subscribers on that downstream RF interface (subject to any administrative controls).
- If the CMTS performs network-layer forwarding of multicast packets, it MUST support Active IGMP mode.
- If link-layer forwarding of multicast packets is used, the CMTS SHOULD support Passive IGMP mode and MAY support Active IGMP mode.

3.3.1.3 CM Rules

The CM MUST support IGMP with the cable-specific rules specified in this section.

The CM MUST implement the passive IGMP mode. Additionally, the CM MAY implement the active IGMP mode. If the CM implements the active IGMP mode, the CM MUST support a capability to switch between modes.

Multicast Forwarding Requirements

The following requirements apply to both passive and active modes of IGMP operations:

- The CM MUST NOT forward Membership Queries from its CPE interface to its RF interface.

- The CM MUST NOT forward Membership Reports or IGMP v2 Leaves received on its RF interface to its CPE interface.
- The CM MUST NOT forward multicast traffic from its RF interface to its CPE interface unless a device on its CPE interface is a member of that IP multicast group.
- The CM MUST forward multicast traffic from its CPE interface to its RF interface unless administratively (via configuration or other mechanism) prohibited.
- As a result of receiving a Membership Report on its CPE interface, the CM MUST begin forwarding traffic for the appropriate IP multicast group. The CM MUST stop forwarding multicast traffic from the RF to the CPE side whenever the CM has not received a Membership Report from the CPE side for more than the Membership Interval, which is (2 * MQI) + QRI, where MQI is the Membership Query Interval and QRI is the Query Response Interval.
- The CM MAY stop forwarding traffic from the RF to the CPE side for a particular multicast group prior to the expiration of the Membership Interval (see above) if it can determine (for example, via an IGMP LEAVE message and the appropriate protocol exchange) that there are no CPE devices subscribed to that particular group.

The following requirements apply only when the CM is operating in passive IGMP mode:

- The CM MUST forward traffic for the ALL-HOSTS multicast group from its RF interface to its CPE interface unless administratively prohibited. The CPE MUST always be considered a member of this group. In particular, the CM MUST forward ALL-HOSTS Group Queries that pass permit filters on its RF interface to its CPE interface.
- Upon receiving a Membership Report on its CPE interface, the CM MUST start a random timer between 0 and 3 seconds. During this time period, the CM MUST discard any additional Membership Reports received in its CPE interface for the associated multicast group. If the CM receives a Membership Report on its HFC interface for the associated multicast group, the CM MUST discard the Membership Report received on its CPE interface. If the random timer expires without the reception of a Membership Report on the CMs HFC interface, the CM MUST transmit the Membership Report received on its CPE interface.

The following requirements apply only when the CM is operating in active IGMP mode:

- The CM MUST implement the Host portion of the IGMP v2 protocol [RFC-2236] on its RF interface for CPEs with active groups and MUST NOT act as a Querier on its RF interface.
- The CM MUST act as an IGMPv2 Querier on its CPE interface.
- If the CM has received a Membership Report on its downstream RF interface for groups active on the CMs CPE interface within the Query Response Interval, it MUST suppress transmission on its upstream RF interface of such Membership Reports.
- The CM MUST suppress all subsequent Membership Reports for this group until such time as the CM receives a Membership Query (General or Specific to this Group) on its RF interface or a IGMPv2 Leave is received for this group from the CPE interface.
- The CM MUST treat Unsolicited Membership Reports (IGMP JOINs) from its CPE interface as a response to a Membership Query received on its RF interface. Upon receipt of this unsolicited JOIN from its CPE interface, the CM MUST start a random timer according to the Host State Diagram, specified in [RFC-2236], and MUST use a Query Response Interval of 3 seconds. As specified above, if the CM receives a Membership Report on its RF interface for this group during this random time period, it MUST suppress transmission of this Join on its upstream RF interface.

Note: Nothing in this section would prohibit the CM from being specifically configured to not forward certain multicast traffic as a matter of network policy.

3.4 Above the Network Layer

The 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]), MUST be supported for network management.
- TFTP (Trivial File Transfer Protocol, [RFC-1350]), a file transfer protocol, MUST be supported for downloading operational software and configuration information, as modified by TFTP Timeout Interval and Transfer Size Options [RFC-2349].
- DHCP (Dynamic Host Configuration Protocol, [RFC-2131]), a framework for passing configuration information to hosts on a TCP/IP network, MUST be supported.
- Time of Day Protocol [RFC-868], MUST be supported to obtain the time of day.

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 messages received on the CMCI port. DHCP, TFTP, and ToD server messages include: DHCPOFFER, 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/IEC10038]. 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 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 destinations match 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 MAC Service Definition

The MAC sublayer service definition is in Appendix E.

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

The Physical Media Dependent sublayer is defined in Section 4 of this document.

3.6.2.1 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

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.

This section applies to the first technology option referred to in Section 1.1. 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.

4.2 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 µsec 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 band 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 4.2.7. Maximum timing error and guard band may vary with CMTSes from different vendors. The term guard time is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.

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 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 Section 4.2.7, Section 4.2.11.1, and Appendix 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 G.

The symbol rate for each upstream channel is defined in an Upstream Channel Descriptor (UCD) MAC message. All CMs using that upstream channel MUST use the defined symbol rate for upstream transmissions.

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, I_1 is the MSB of the symbol map, Q_1 is the LSB for QPSK, and Q_0 is the LSB for 16QAM. Q_1 and I_0 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.

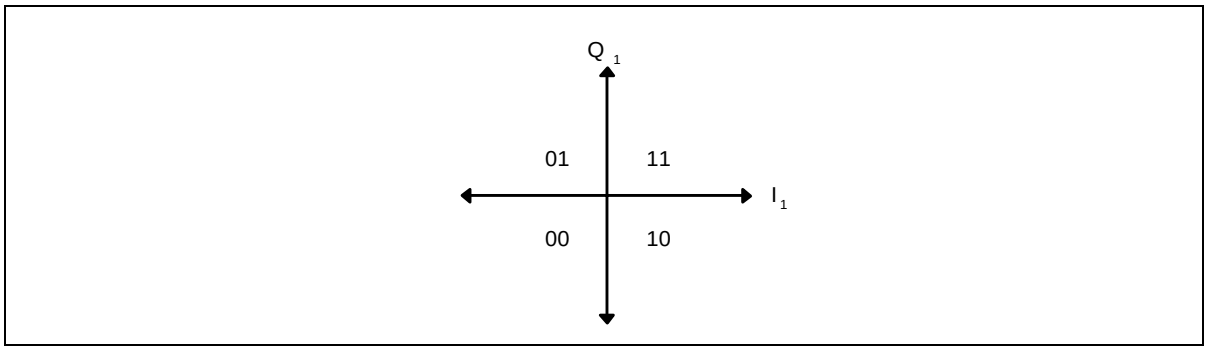


Figure 4-1. QPSK Symbol Mapping

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

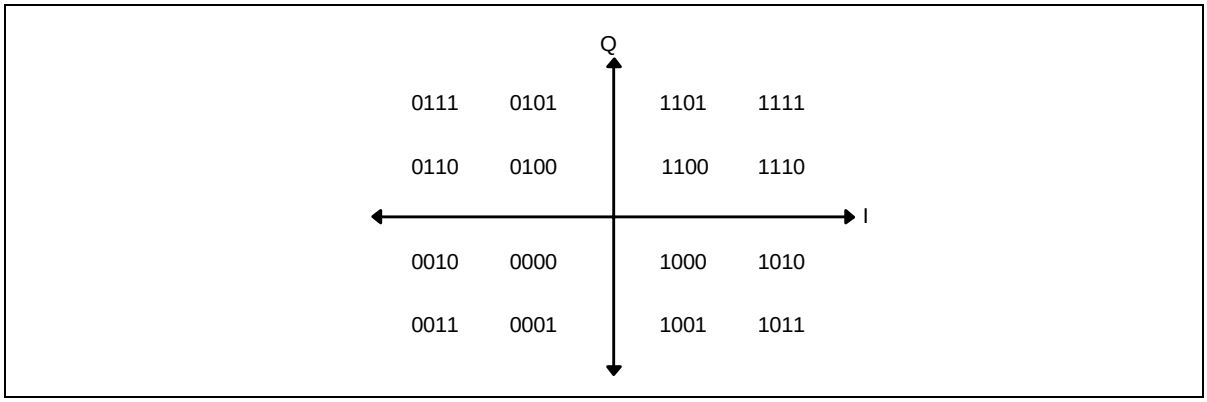


Figure 4-2. 16QAM Gray-Coded Symbol Mapping

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

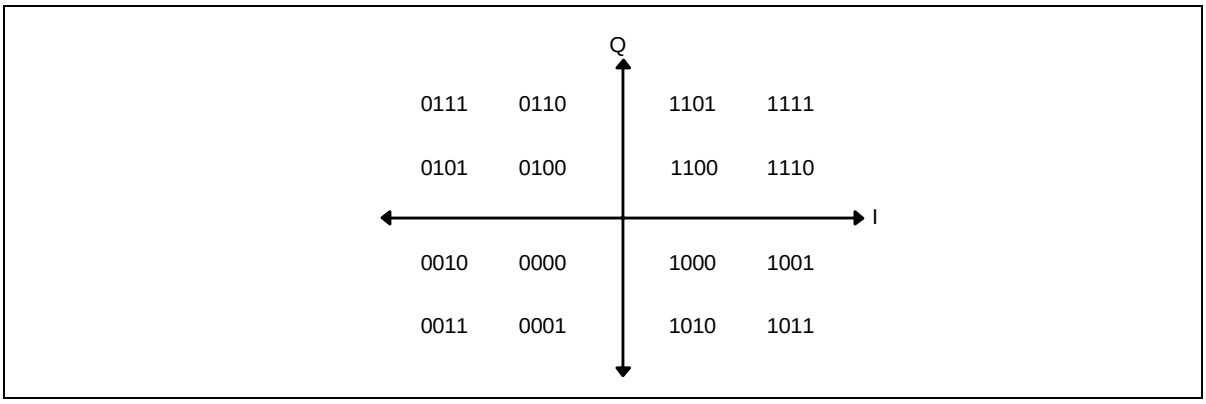


Figure 4-3. 16QAM Differential-Coded Symbol Mapping

If differential quadrant encoding is enabled, then the currently-transmitted symbol quadrant is derived from the previously transmitted symbol quadrant and the current input bits via Table 4-2. If differential quadrant encoding is enabled, the upstream PMD sublayer MUST apply these differential encoding rules to all transmitted symbols (including those that carry preamble bits).

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)[1]
160	200
320	400
640	800
1,280	1,600
2,560	3,200

Notes to Table 4-1:

1. 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 $s(t) = I(t)*\cos(\omega t) - Q(t)*\sin(\omega t)$, 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:

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

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 minimum uncoded word size MUST be 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 4.2.11.1.2 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 $x^{15} + x^{14} + 1$.

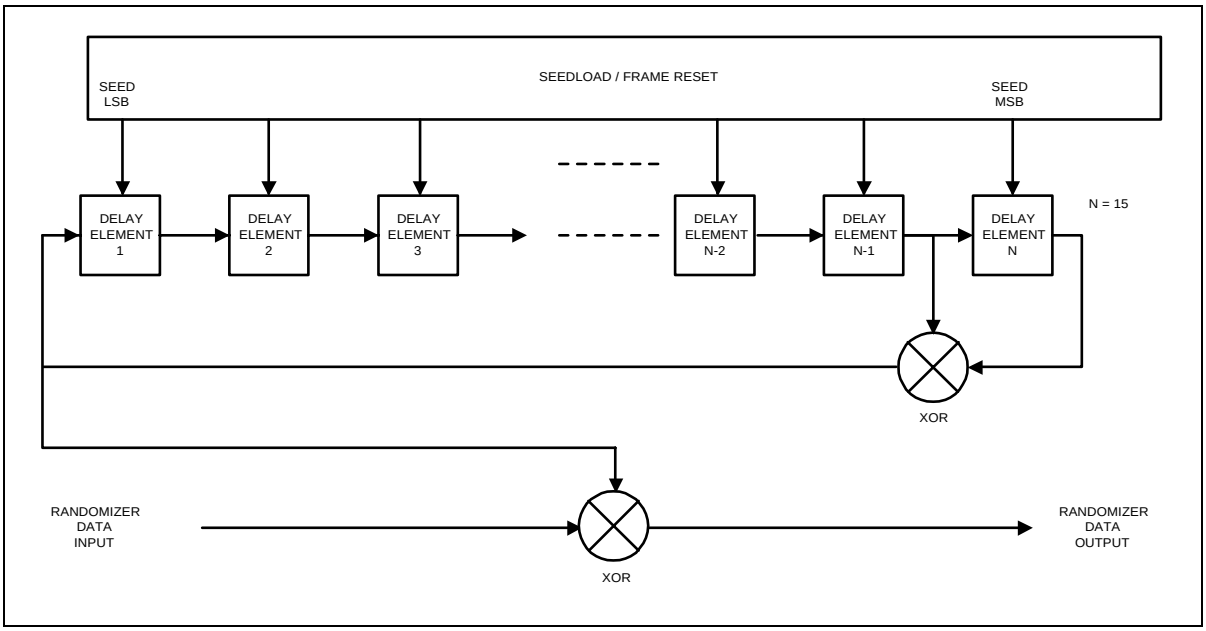


Figure 4-4. Scrambler Structure

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 4-9), and is I_1 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 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 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

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.

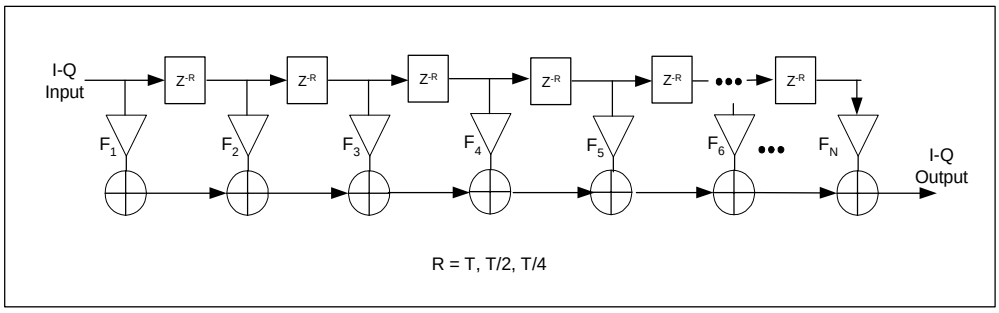


Figure 4-5. Transmit Pre-Equalizer Structure

The RNG-RSP MAC message, (see Section 6.3.6.1) 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 6.3.8.1.1 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 though 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 NOT change its commanded output transmit power due to a gain or loss of the new coefficients. The actual output transmit power is subject to the power accuracy requirements defined in Section 4.2.9.1. 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 6.3.3 Table 6-18; 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 6.3.3 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 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 6.3.4).

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).

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)	0 to 10 (0 implies no FEC. The number of codeword parity bytes is 2*T)
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)[1]	0 to 255
Guard Time	4 to 255 symbols
Last Codeword Length	Fixed, shortened
Scrambler On/Off	On/Off

Notes to Table 4-4:

1. 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 4-5. User Unique Burst Parameters

User Unique Parameter	Configuration Settings
Power Level[1]	+8 to +55 dBmV (16QAM) +8 to +58 dBmV (QPSK) 1-dB steps
Offset Frequency[1]	Range = ±32 kHz; increment = 1 Hz; implement within ±10 Hz
Ranging Offset	0 to (2^{16} - 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[1]	Up to 64 coefficients; 4 bytes per coefficient:2 real and 2 complex

Notes to Table 4-5:

1. 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. It is an advancement equal to roughly the round-trip delay of the CM from the CMTS, and is needed to synchronize upstream transmissions in the TDMA scheme. 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 µsec/64 = 0.09765625 µsec = 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 ±0.25 µsec plus ±1/2 symbol owing to resolution. The accuracy of CM burst timing of ±0.25 µsec plus ±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 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. 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 transmittter 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. To assure that this does not happen, the CMTS MUST allow extra guard time between bursts that is at least equal to the amount of negative ranging offset.

To provide backward interoperability with DOCSIS 1.0 and 1.1 equipment, when making a symbol rate change the CM MUST employ the following timing offsets when changing symbol rates. The offsets in the table correspond to the contribution of DOCSIS 1.0 and 1.1 legacy upstream receivers to changes in latency when making symbol rate changes. The timing offset to apply is the difference between the entry in the table corresponding to the new symbol rate and the entry corresponding to the original symbol rate. The offsets are referenced to the center of the first symbol in the burst, which is the reference point for burst timing as stated in Section 4.2.8. Specification of these offsets is needed so that CMs apply uniform adjustments to their ranging offsets and so that CMTSes can appropriately handle CMs that apply these offsets when making symbol rate changes.

Symbol Rate	Timing Offset (in units of 1/64 time ticks referenced to 2.56 Msps)
2.56 Msps	0 (reference)
1.28 Msps	24
0.64 Msps	72
0.32 Msps	168
0.16 Msps	360

As an example, suppose a CM is on an upstream channel operating at 1.28 Msps. Now, suppose the UCD message from the CMTS changes the symbol rate of the channel to 0.32 Msps. The CM applies an additional timing offset of 168 - 24 = 144 to its ranging offset to compensate for this symbol rate change. The value of 144 is positive, and thus, the CM will add to its ranging offset so that it effectively transmits earlier by 144 units of 1/64 time ticks.

Furthermore, in changing symbol rates, if a CM has its own contribution to a change in latency, the CM MUST also compensate for this CM-specific latency difference. This is in addition to the offset applied from the values in the table above, which result from legacy CMTS upstream receiver contributions to changes in latency. The requirements for CM burst timing accuracy found earlier in this section, referenced to the symbol rate that is the lower of the original and the new symbol rate, apply after the symbol rate change with the required timing offsets above considered.

A CMTS that does not apply the same internal physical delay offsets as the legacy DOCSIS upstream CMTS receiver implementation is capable of receiving a CM burst after a symbol rate change in any of the following ways but is not limited necessarily to only these ways:

a) The CMTS may implement the internal physical delay offset, as specified in the above table.

b) The CMTS may implement an internal timing compensation based on the expected offset in the above table.

c) The CMTS may increase the guard time.

d) The CMTS may send an unsolicited RNG-RSP to each CM to adjust the delay offset. As discussed in Section 6.3.6, the CM is expected to be capable of adjusting its timing offset at any time with the accuracy specified within this section.

If Channel Frequency is to be changed, then 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 Section 4.2.10.5 and Section 4.2.10.6 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 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 16 QAM).

4.2.8 Burst Timing Convention

Figure 4-6 illustrates the nominal burst timing.

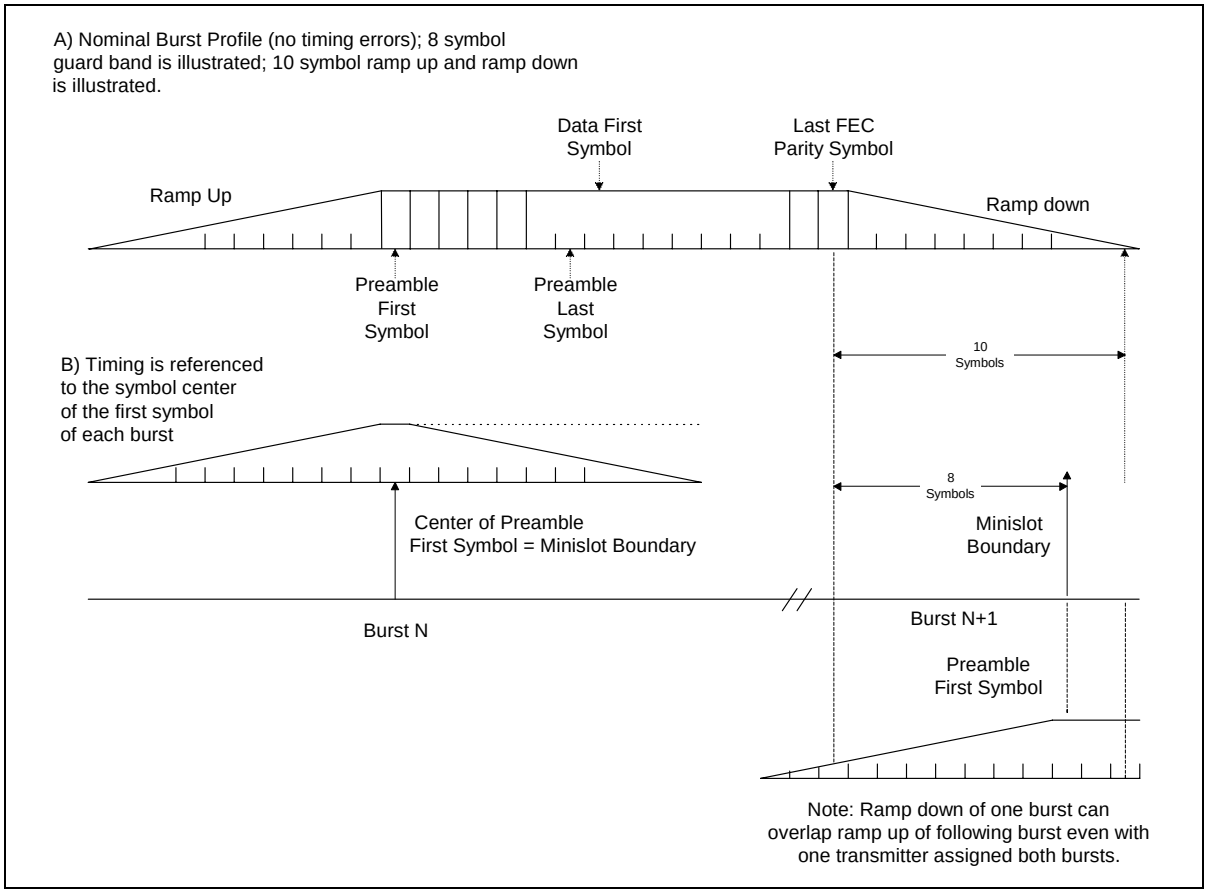


Figure 4-6. Nominal Burst Timing

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.

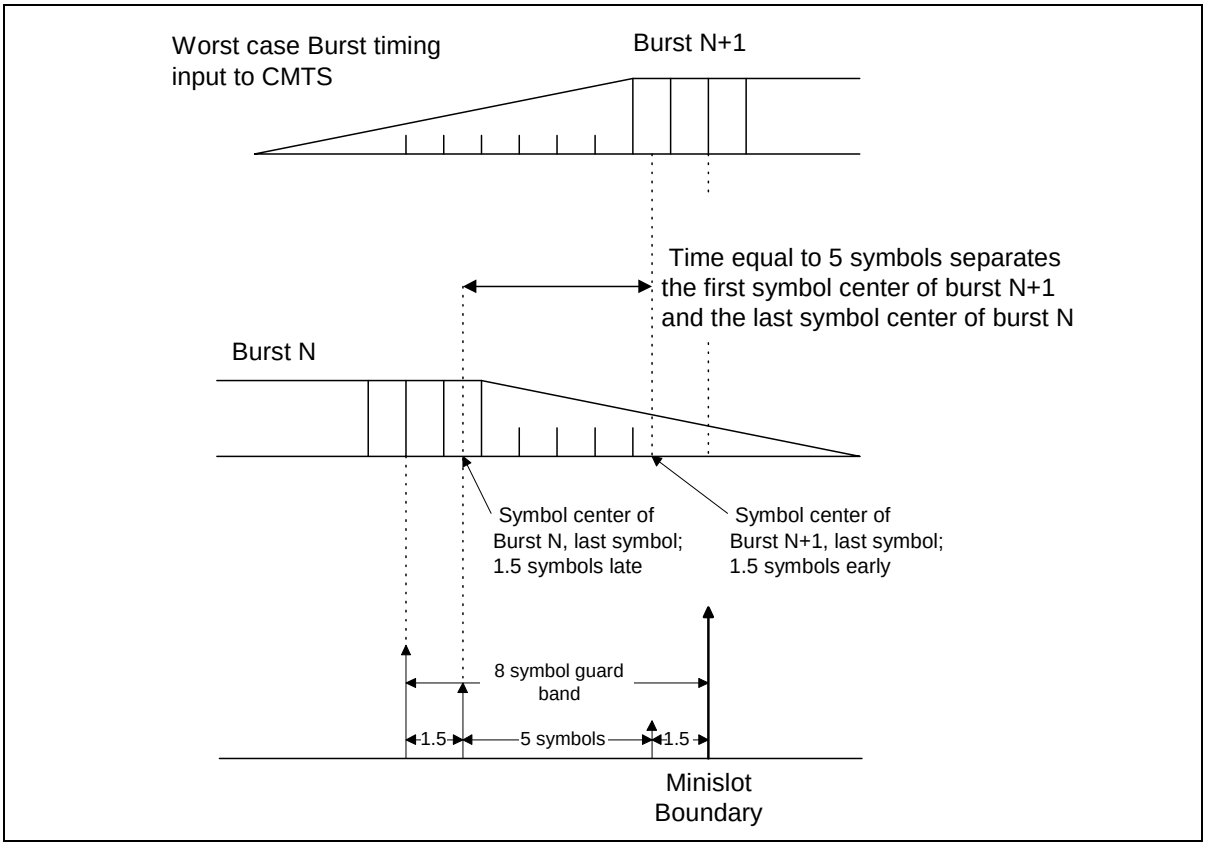


Figure 4-7. Worst-Case Burst Timing

At a symbol rate of Rs, symbols occur at a rate of one each Ts = 1/Rs seconds. Ramp Up and Ramp Down are the spread of a symbol in the time domain beyond Ts duration owing to the symbol-shaping filter. If only one symbol were transmitted, its duration would be longer than Ts due to the shaping filter impulse response being longer than Ts. The spread of the first and last symbols of a burst transmission effectively extends the duration of the burst to longer than N * Ts, 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 9.2.4 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 ± 2dB, and the step size accuracy ± 0.4dB, 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 ± 1.4dB. 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 three 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. Carrier-related spurious emissions include all products whose frequency is a function of the carrier frequency of the upstream transmission, such as but not limited to carrier harmonics.

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)[1]		
42 to 54 MHz 54 to 60 MHz 60 to 88 MHz 88- to 860 MHz	max(-40 dBc, -26 dBmV) -35 dBmV -40 dBmV -45 dBmV	-26 dBmV -40 dBmV -40 dBmV max(-45 dBmV, -40 dBc)[2]
CM Discrete Spurious Emissions Limits:[1]		
42 to 54 MHz 54 to 88 MHz 88 to 860 MHz	-max(-50 dBc, -36 dBmV) -50 dBmV -50 dBmV	-36 dBmV -50 dBmV -50 dBmV

Notes to Table 4-6:

1. These specification limits exclude a single discrete spur related to the tuned received channel; this single discrete spur MUST be no greater than -40 dBmV.
2. “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 mChannel Spurious Emissions Relative to the Transmitted Burst Power Level

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	2560 Ksym/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 Relative to the Transmitted Burst Power Level

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_o), 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.

NOTE: The CM MUST be capable of achieving a cluster SNR of at least 27 dB in the presence of the channel micro-reflections defined in Table 2-2. Since the table does not bound echo delay for the -30 dBc case, for testing purposes it is assumed that the time span of the echo at this magnitude is less than or equal to 1.5 us.

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 in Table 4-9.

Table 4-9. Filter Amplitude Distortion

Frequency	Amplitude Range	
	low	high
$f_c - 5R_s/8$	-	-30dB
$f_c - R_s/2$	-3.5dB	-2.5dB
$f_c - 3R_s/8$ to $f_c - R_s/4$	-0.5dB	+0.3dB
$f_c - R_s/4$ to $f_c + R_s/4$	-0.3dB	+0.3dB
$f_c + R_s/4$ to $f_c + 3R_s/8$	-0.5dB	+0.3dB
$f_c + R_s/2$	-3.5dB	-2.5dB
$f_c + 5R_s/8$	-	-30dB

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 ± 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 ± 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 4.2.11.1.

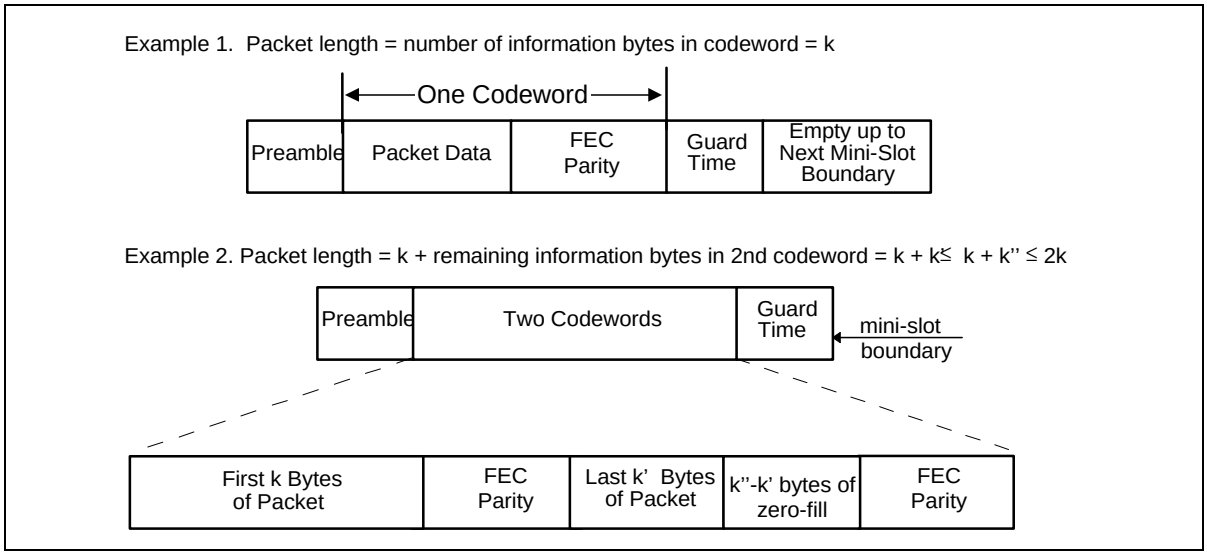


Figure 4-8. Example Frame Structures with Flexible Burst Length Mode

4.2.11.1 Codeword Length

When FEC is enabled, the CM operates in either fix-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.

Shortened last codeword mode MUST NOT be used for Initial Maintenance.

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 4-8, 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 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 MUST 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" - k' zero-fill bytes are necessary to fill out a mini-slot allocation with 16 <= k" <= k and k' <= 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 MUST 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, then 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.

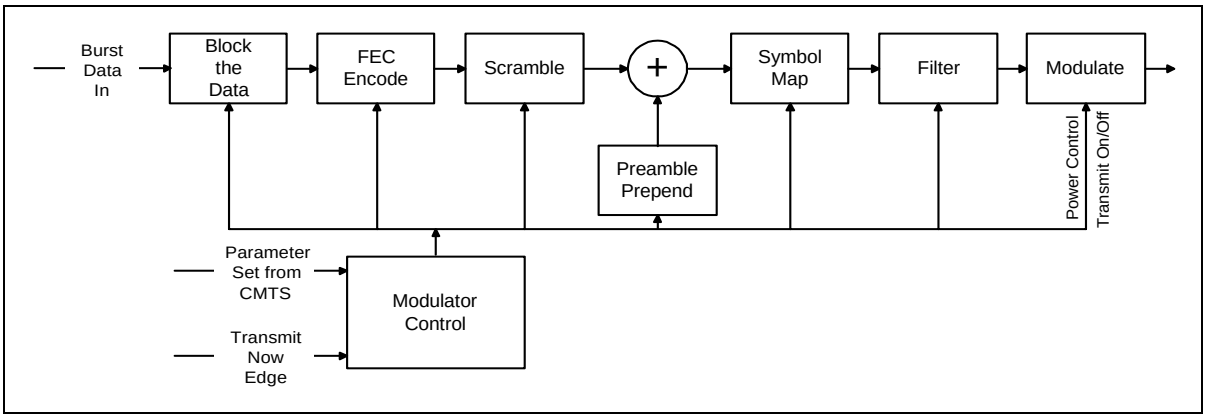


Figure 4-9. Signal-Processing Sequence

	Packet Stream Input
	⇓
Block the Data	Separate Packet into Information Blocks(=data bytes in one codeword)
	⇓
FEC Encode	FEC Encode (Reed-Solomon)each Information Block, using shortened codeword for last block if needed.FEC can be turned off.
	⇓
Scramble	Scramble (see Section 4.2.4)
	⇓
Preamble Prepend	Preamble prepend Symbols
	⇓
Symbol Map	Symbol Map the Data Stream into Modulator Symbols
	⇓
Filter	Filter Symbol Stream for Spectral Shaping
	⇓
Modulate	Modulate at Precise Times(QPSK; 16 QAM)
	⇓
	Output RF Waveform Bursts

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 4-10.

Table 4-10. 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 4-11.

Table 4-11. 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	F connector per [ISO-169-24] (common with the input)

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-T 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 Section 1.1 (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 4-12. The table contains a subset of interleaver modes found in [ITU-T J.83-B].

Table 4-12. Interleaver Characteristics

I (Number of Taps)	J (Increment)	Burst Protection 64QAM/256QAM	Latency 64QAM/256QAM
8	16	5.9 µsec/4.1 µsec	0.22 msec/0.15 msec
16	8	12 µsec/8.2 µsec	0.48 msec/0.33 msec
32	4	24 µsec/16 µsec	0.98 msec/0.68 msec
64	2	47 µsec/33 µsec	2.0 msec/1.4 msec
128	1	95 µsec/66 µsec	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-T 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 [EIA-S542]. 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 4-13.

Table 4-13. CMTS Output

Parameter	Value
Center Frequency (fc)	91 to 857 MHz ±30 kHz[1]
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 (fc ± 3 MHz)	< -57dBc
Inband Spurious and Noise (fc ± 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 (fc ± 3.0 MHz) to (fc ± 3.75 MHz)	< -58 dBc in 750 kHz
Adjacent channel (fc ± 3.75 MHz) to (fc ± 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 (fc ± 9 MHz) to (fc ± 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	F connector per [ISO-169-24]]

Notes to Table 4-13:

1. ±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 be able to locate and accept RF modulated signals located within channels defined in [EIA-S542] for Harmonic Related Carrier (HRC), Incremental Related Carrier (IRC), and Standard (STD) North American frequency plans. Operation below a center frequency of 91 MHz is not required. The signals will have the characteristics defined in Table 4-14.

Table 4-14. Electrical Input to CM

Parameter	Value
Center Frequency	91 to 857 MHz ±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	F connector per [ISO-169-24] (common with the output)

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 a digital signal at 0 dBc in the adjacent channels.

Performance as described in Section 4.3.6.1.1 MUST be met with an 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 a 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.

Input Receive Signal Level	E_S/N_o
-6 dBmV to +15dBmV	30dB or greater
Less than -6dBmV down to -15dBmV	33dB 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 an analog or a 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 an analog or a 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 an 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 a digital signal at +10 dBc in the adjacent channels.

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:

$$| (N2-N1)/f_{CMTS} - (T2-T1) | < 500x10^{-9}$$

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 500 nsec 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 4.2.7 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 accuracy of ≤ ±5 ppm, drift rate ≤ 10^{-8} per second, and edge jitter of ≤ 10 nsec peak-to-peak (±5 nsec) over a temperature range of 0 to 40 degrees C up to ten years from date of manufacture.[1] [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 segments' 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 one 10,240,000 cycle segment, ≤ 30 nsec; adjacent 102,400,000 segments, ≤ 120 nsec. The CMTS Master Clock MUST meet such test limits in 99% or more measurements.]

1. This specification MAY also be met by synchronizing the CMTS Master Clock oscillator to an external frequency reference source. If this approach is used, the internal CMTS Master Clock MUST have frequency accuracy of ±20ppm over a temperature range of 0 to 40 degrees C up to 10 years from date of manufacture when no frequency reference source is connected. The drift rate and edge jitter MUST be as specified above.

5 Downstream Transmission Convergence Sublayer

5.1 Introduction

This section applies to the first technology option referred to in Section 1.1(Scope). For the second option, refer to Appendix N.

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).

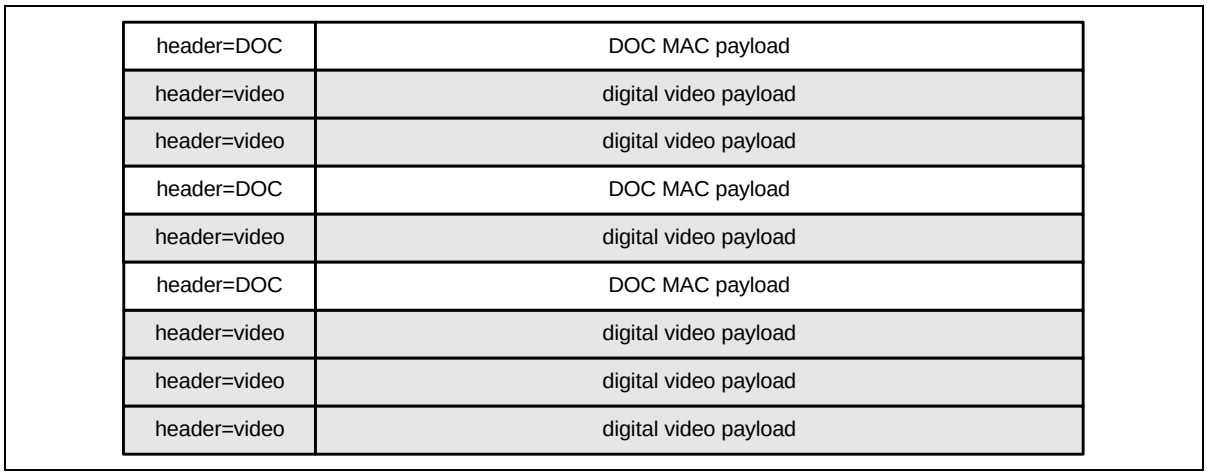


Figure 5-1. Example of Interleaving MPEG Packets in Downstream

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.

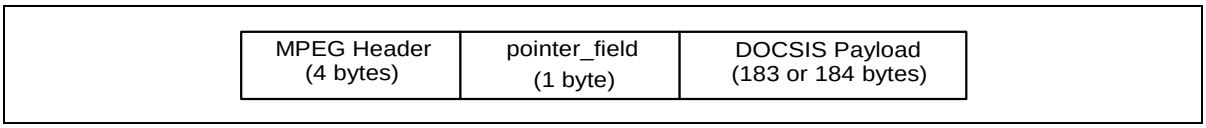


Figure 5-2. Format of an MPEG Packet

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
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	13	DOCSIS 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 DOCSIS PID
continuity_counter	4	cyclic counter within this PID

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.

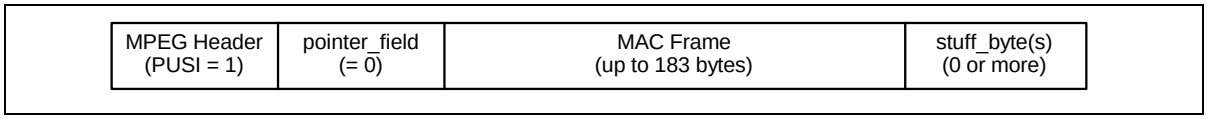


Figure 5-3. Packet Format Where a MAC Frame Immediately Follows the pointer_field

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.

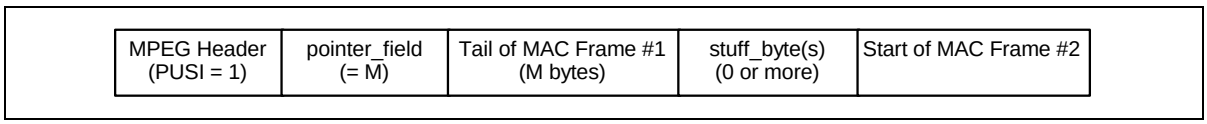


Figure 5-4. Packet Format with MAC Frame Preceded by Stuffing Bytes

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.

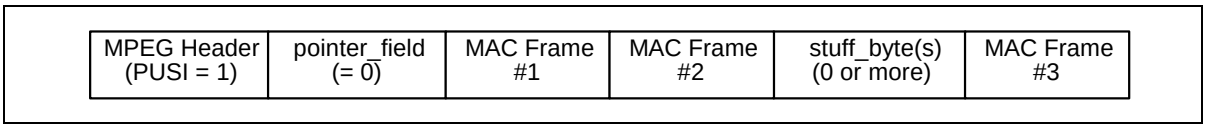


Figure 5-5. Packet Format Showing Multiple MAC Frames in a Single Packet

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.

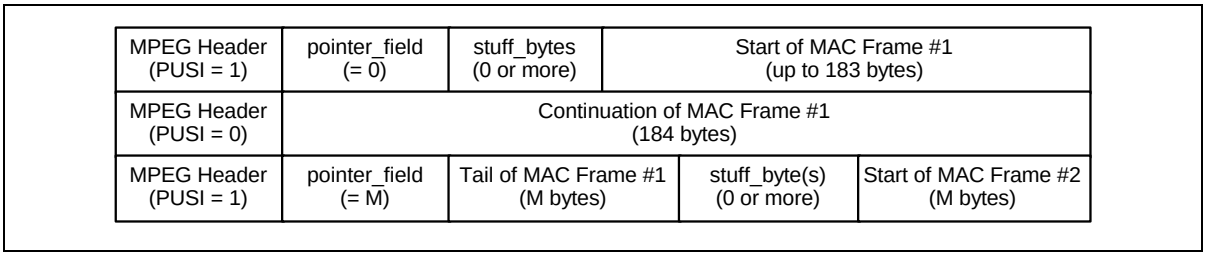


Figure 5-6. Packet Format Where a MAC Frame Spans Multiple Packets

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 and Section 7.3).

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 1.1 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
- Quality-of-service including:
    - Support for Bandwidth and Latency Guarantees
    - Packet Classification
    - Dynamic Service Establishment
- 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. (Refer to Section 3.2)

6.1.2.3 Service Flows

The concept of Service Flows is central to the operation of the MAC protocol. Service Flows provide a mechanism for upstream and downstream Quality of Service management. In particular, they are integral to bandwidth allocation.

A Service Flow ID defines a particular unidirectional mapping between a CM and the CMTS. Active Upstream Service Flow IDs also have associated Service IDs or SIDs. Upstream bandwidth is allocated to SIDs, and hence

to CMs, by the CMTS. Service IDs provide the mechanism by which upstream Quality of Service is implemented.

The CMTS MAY assign one or more Service Flow IDs (SFIDs) to each CM, corresponding to the Service Flows required by the CM. This mapping can be negotiated between the CMTS and the CM during CM registration or via dynamic service establishment (refer to Section 9.4).

In a basic CM implementation, two Service Flows (one upstream, one downstream) could be used, for example, to offer best-effort IP service. However, the Service Flow concept allows for more complex CMs to be developed with support for multiple service classes while supporting interoperability with more basic modems. With these more complex modems, it is possible that certain Service Flows will be configured in such a way that they cannot carry all types of traffic. That is, they may have a maximum packet size limitation or be restricted to small fixed size unsolicited grants. Furthermore it might not be appropriate to send other kinds of data on Service Flows that are being used for Constant Bit Rate (CBR)-type applications.

Even in these complex modems, it is necessary to be able to send certain upstream packets needed for MAC management, SNMP management, key management, etc. For the network to function properly, all CMs MUST support at least one upstream and one downstream Service Flow. These Service Flows MUST always be provisioned to allow the CM to request and to send the largest possible unconcatenated MAC frame (refer to Section 6.2.2). These Service Flows are referred to as the upstream and downstream Primary Service Flows. The SID assigned to the upstream Primary Service Flow is referred to as the Primary SID.

The Primary SID MUST always be assigned to the first provisioned upstream Service Flow during the registration process (which may or may not be the same temporary SID used for the registration process). The Primary Service Flows MUST be immediately activated at registration time. The Primary SID MUST always be used for station maintenance after registration. The Primary Service Flows MAY be used for traffic. All unicast Service Flows MUST use the security association defined for the Primary Service Flow. (Refer to [DOCSIS8])

All Service Flow IDs are unique within a single MAC-sublayer domain. The mapping of a unicast Service Identifier to an active/admitted Service Flow MUST be unique within a single MAC-sublayer domain. The length of the Service Flow ID is 32 bits. 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 6.25µs increments, *i.e.*, 2, 4, 8, 16, 32, 64, or 128 times 6.25µs. The relationship between mini-slots, bytes, and time ticks is described further in Section 7.3.4. The usage code values are defined in Table 6-20 and allowed use is defined in Section 6.3. The binding of these values to physical-layer parameters is defined in Table 6-18.

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 a variable-length data PDU. The variable-length PDU includes a pair of 48-bit addresses, data, and a CRC. In

special cases, the MAC Header may encapsulate multiple MAC frames (see Section 6.2.5.5) into a single MAC frame.

6.1.3 Future Use

A number of fields are defined as being "for future use" or Reserved in the various MAC frames described in this document. These fields MUST NOT be interpreted or used in any manner by this version (1.1) 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.

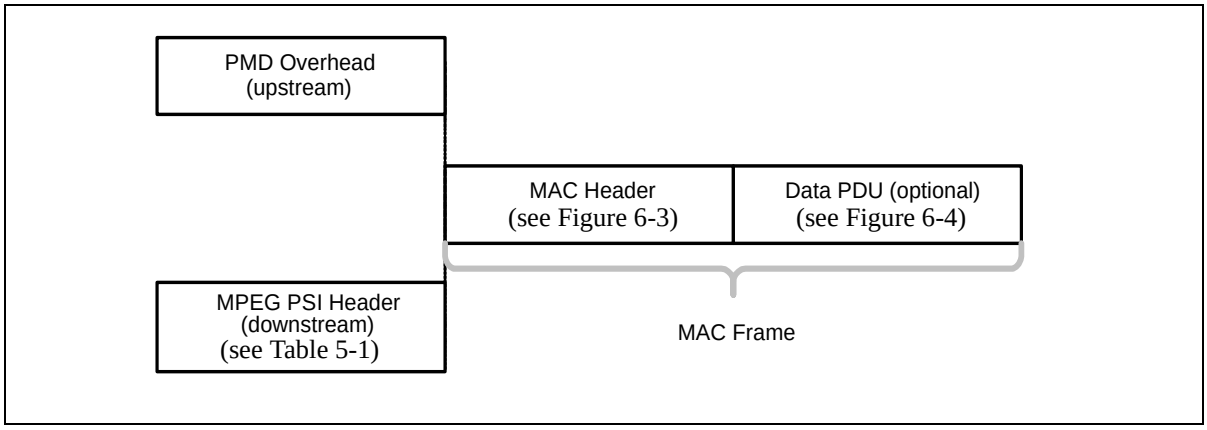


Figure 6-1. Generic MAC Frame Format

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 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 (refer to Section 7.1).

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.

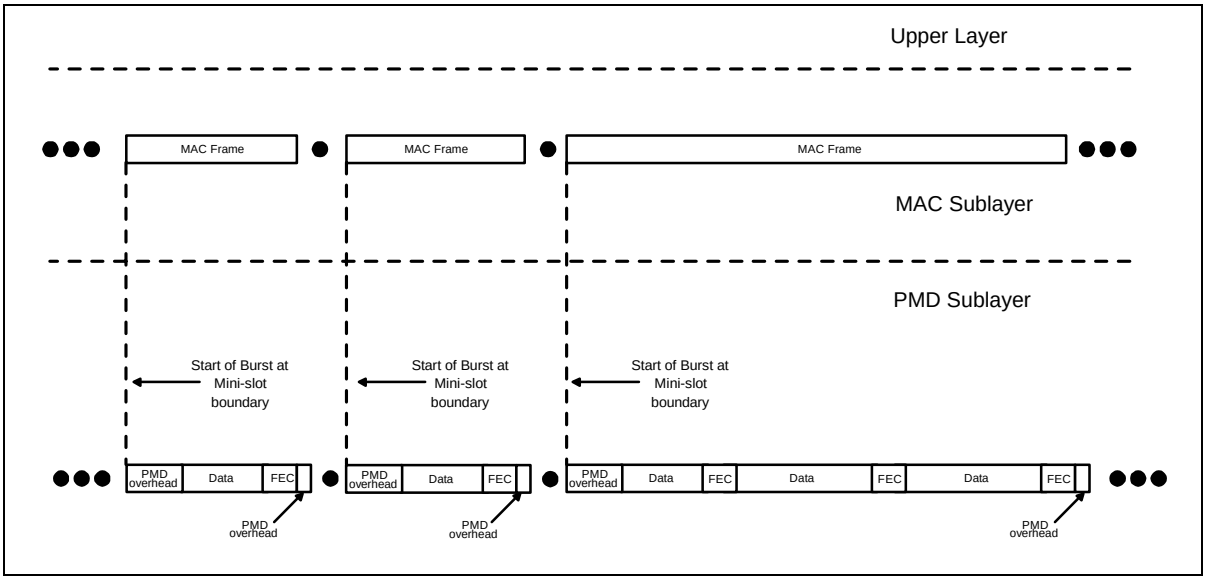


Figure 6-2. Upstream MAC/PMD Convergence

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[ISO8802-3]. This is often called bit-little-endian order.[1]

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 is sometimes called byte-big-endian order.

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.

6.2.1.3.1 Representing Negative Numbers

Signed integer values MUST 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 are unordered lists of TLV-tuples. Some TLV's are nested (see Appendix C). All TLV Length fields, except for EH_LEN (see Section 6.2.6), MUST be greater than zero. Unless otherwise specified, Type is one byte and Length is one byte.

Using this encoding, new parameters may be added which some devices cannot interpret. A CM or CMTS which does not recognize a parameter type MUST skip over this parameter and MUST NOT treat the event as an error condition.

6.2.1.4 MAC Header Format

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

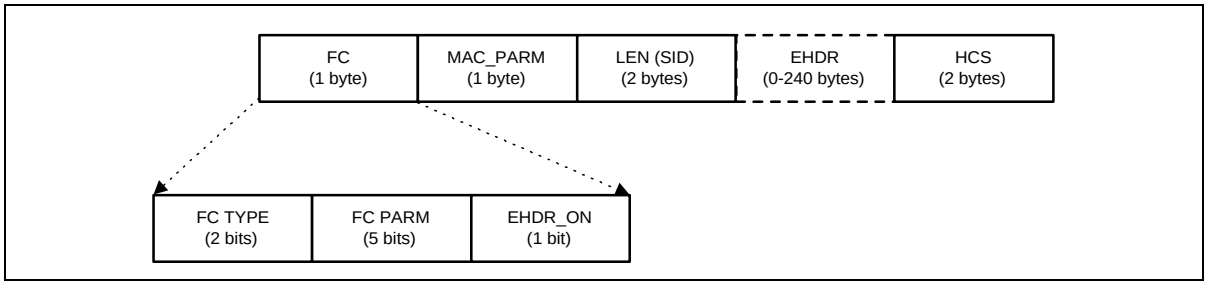


Figure 6-3. MAC Header Format

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

1. 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.

bytes of MAC control; an OPTIONAL Extended Header field (EHDR); plus a Header Check Sequence (HCS) to ensure the integrity of the MAC Header.

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 6-10) used for MAC frame count else (for Requests only) indicates the number of mini-slots 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, then 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, then 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, then the MAC_PARM field represents the number of MAC frames (CNT) in the concatenation (see Section 6.2.5.5). If this is a Request MAC Header (REQ) (see Section 6.2.5.3), then 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 as well as frame fragmentation, 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-Specific Frame 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. Normally, the Packet PDU MUST be passed across the network in its entirety, including its original CRC.[2] A unique Packet

[2] The one exception is the case of Payload Header Suppression. In this case, all bytes except those suppressed MUST be passed across the network and the CRC covers only those bytes actually transmitted. (Refer to Section 6.2.6.3.1)

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.

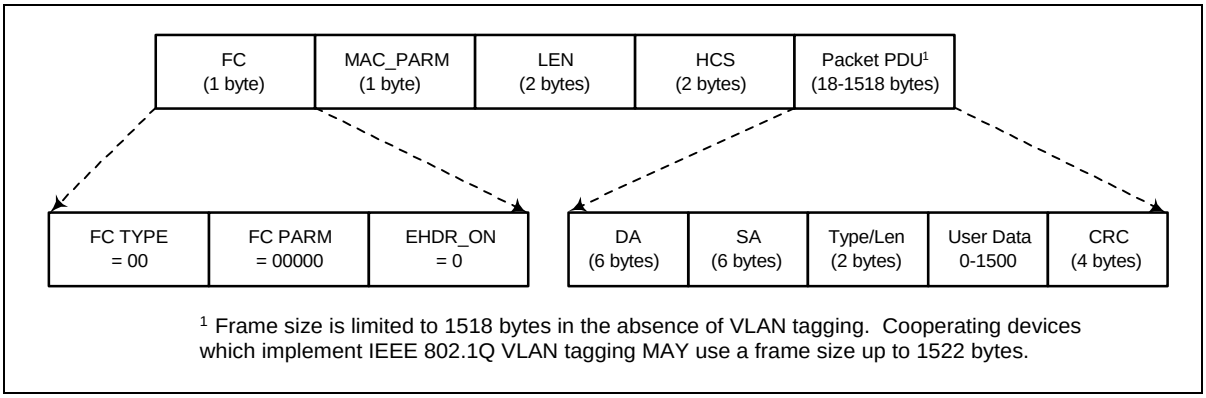


Figure 6-4. Ethernet/802.3 Packet PDU Format

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 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n+x; length of Packet PDU in bytes + length of EHDR	16 bits
EHDR	Extended MAC Header, if present	x (0-240) bytes
HCS	MAC Header Check Sequence	16 bits
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 + x + n bytes

Under certain circumstances (see appendix M) it may be necessary to transmit a packet PDU MAC frame without an actual PDU. This is done so that the extended header can be used to carry certain information about the state of the service flow. This could also happen as a result of PHS (see Section 6.2.6.3.1). Such a frame will have the length field in MAC header set to the length of the extended header and will have no packet data, and therefore no CRC. This can only happen with frames transmitted on the upstream as frames transmitted on the downstream always have at least the DA and SA fields of the packet PDU.

6.2.3 ATM Cell MAC Frames

The FC_TYPE 0x01 is reserved for future definition of ATM Cell MAC Frames. This FC_TYPE field in the MAC Header indicates that an ATM PDU is present. This PDU MUST be silently discarded by MAC

implementations of this version (1.1) of the specification. Compliant version 1.1 implementations MUST use the length field to skip over the ATM PDU.

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 (1.1) of the specification. Compliant version 1.1 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 6-5 and Table 6-4.

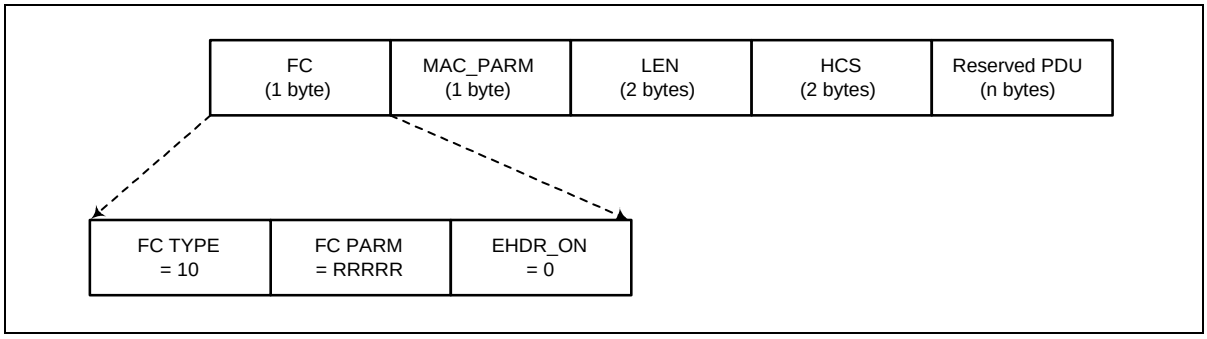


Figure 6-5. Reserved PDU Format

Table 6-4. 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 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n+x; length of Reserved PDU + length of EHDR in bytes	16 bits
EHDR	Extended MAC Header, if present	x (0-240) bytes
HCS	MAC Header Check Sequence	16 bits
User Data	Reserved Data PDU	n bytes
	Length of Reserved PDU MAC frame	6 + x + 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, fragmentation and concatenating multiple MAC frames.

Table 6-5 describes FC_PARM usage within the MAC Specific Header.

FC_PARM	Header/Frame Type
00000	Timing Header
00001	MAC Management Header
00010	Request Frame
00011	Fragmentation Header
11100	Concatenation Header

Table 6-5. MAC-Specific Headers and 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 6-6 and Table 6-6.

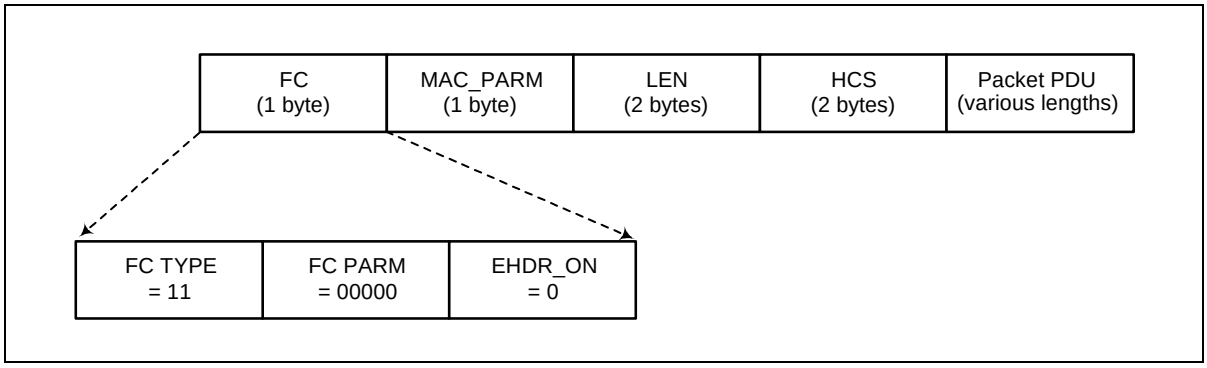


Figure 6-6. Timing MAC Header

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 6-7 and Table 6-7.

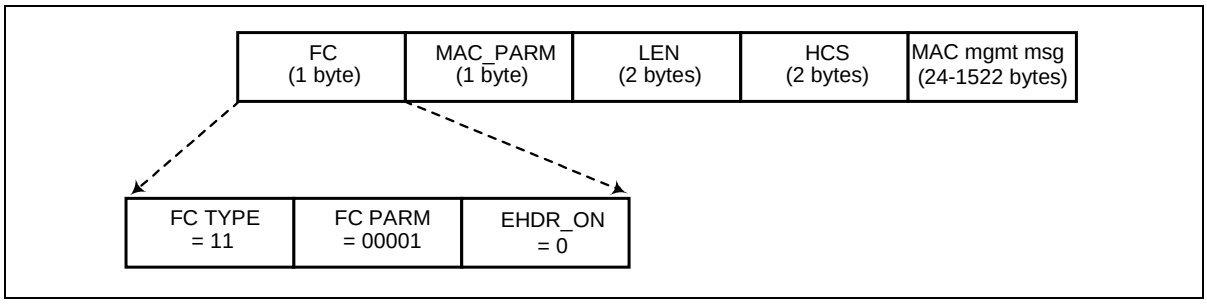


Figure 6-7. Management MAC Header

Table 6-7. MAC Management Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC Specific Header FC_PARM[4:0] = 00001; Management MAC Header EHDR_ON = 0 if there is no extended header, 1 if there is an EHDR	8 bits
MAC_PARM	MAC_PARM = x; MUST be set to zero if there is no EHDR; Otherwise set to length of EHDR	8 bits
LEN	LEN = n+x; length of MAC management message + length of EHDR in bytes	16 bits
EHDR	Extended MAC Header, if present	x (0-240) bytes
HCS	MAC Header Check Sequence	16 bits
Packet Data	MAC management message	n bytes
	Length of Packet MAC frame	6 + x + n bytes

6.2.5.3 Request Frame

The Request Frame 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 Frame. The general format of the Request MUST be as shown in Figure 6-8 and Table 6-8.

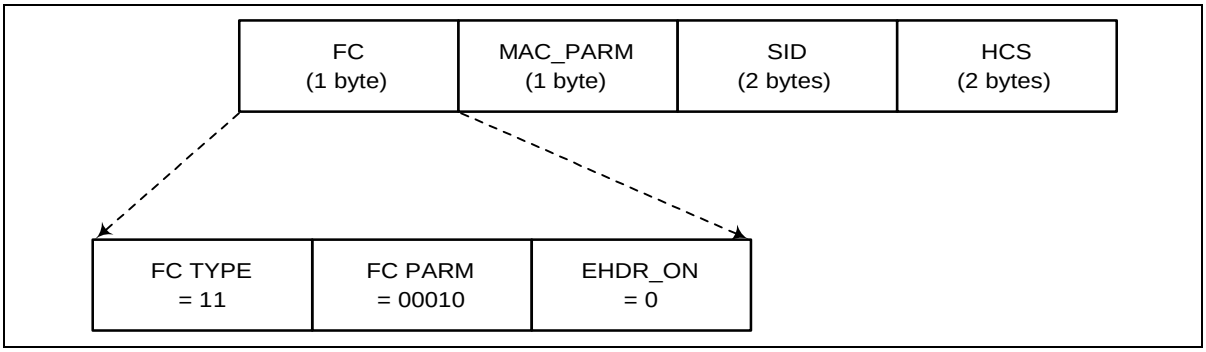


Figure 6-8. Request Frame Format

Table 6-8. Request Frame (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 number of minislots requested	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 Frame 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 Flow within a given CM.

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 Fragmentation Header

The Fragmentation MAC Header provides the basic mechanism to split a larger MAC PDU into smaller pieces that are transmitted individually and then re-assembled at the CMTS. As such, it is only applicable in the upstream. The general format of the Fragmentation MAC Header MUST be as shown in Figure 6-9.

A compliant CM MUST support fragmentation. A compliant CMTS MAY support fragmentation. To decrease the burden on the CMTS and to reduce unnecessary overhead, fragmentation headers MUST NOT be used on unfragmented frames.

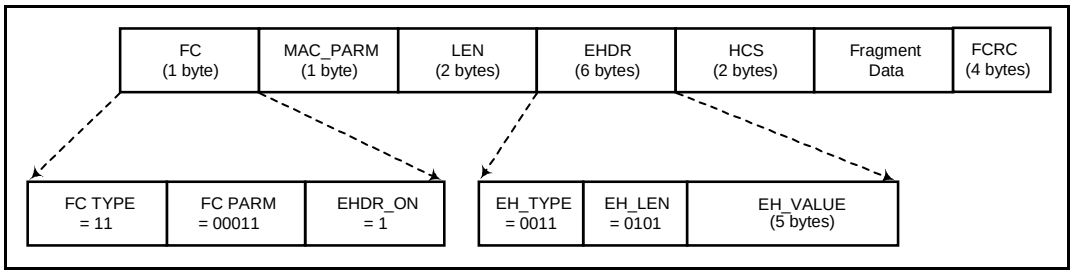


Figure 6-9. Fragmentation MAC Header Format

Table 6-9. Fragmentation MAC Frame (FRAG) Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-Specific Header FC_PARM [4:0] = 00011; Fragmentation MAC Header EHDR_ON = 1; Fragmentation EHDR follows	8 bits
MAC_PARM	ELEN = 6 bytes; length of Fragmentation EHDR	8 bits
LEN	LEN = length of fragment payload + EHDR length + FCRC length	16 bits
EHDR	Refer to Section 6.2.6.2	6 bytes
HCS	MAC Header Check Sequence	2 bytes
Fragment Data	Fragment payload; portion of total MAC PDU being sent	n bytes
FCRC	CRC - 32-bit CRC over Fragment Data payload (as defined in Ethernet/[ISO8802-3])	4 bytes
	Length of a MAC Fragment Frame	16 + n bytes

6.2.5.5 Concatenation Header

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[3] and the Concatenation MAC Header only occur once. Concatenation of multiple MAC frames MUST be as shown in Figure 6-10. Concatenation of multiple MAC frames is the only method by which the CM can transmit more than one MAC frame in a single transmit opportunity.

A compliant CM MUST support concatenation. A compliant CMTS MAY support concatenation. Concatenation only applies to upstream traffic. Concatenation MUST NOT be used on downstream traffic.

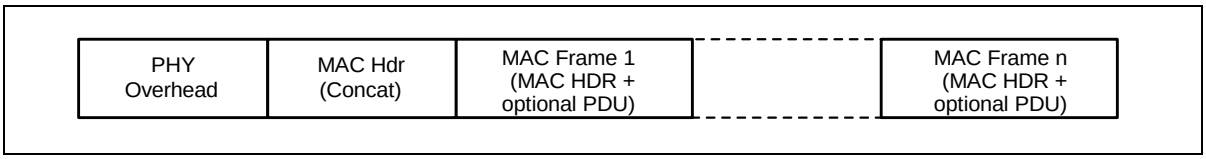


Figure 6-10. Concatenation of Multiple MAC Frames

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 and MAC-specific Frames MAY be mixed together. However, all frames in a concatenation MUST be assigned to the same Service Flow. If the CMTS supports concatenation, it

3. This includes the preamble, guard time, and possibly zero-fill bytes in the last codeword. The FEC overhead recurs for each codeword.

MUST support concatenations containing multiple frame types, including both Packet and MAC-specific frames.

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 6-10.

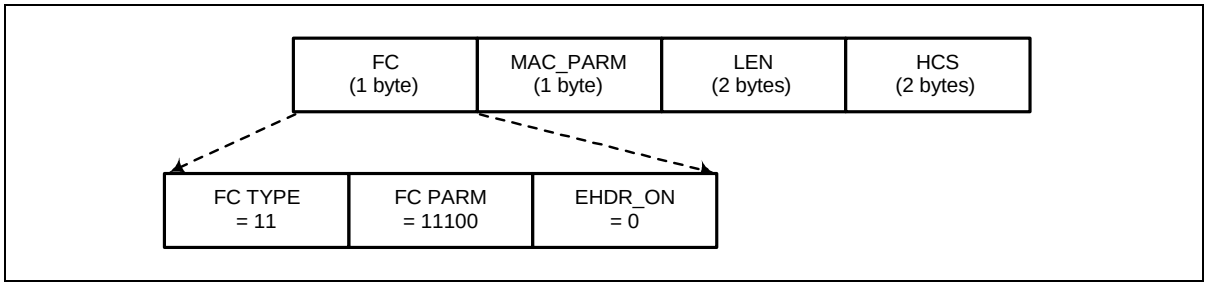


Figure 6-11. Concatenation MAC Header Format

Table 6-10. 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, then it MUST indicate the total count of MAC Frames (CNT) in this concatenantion 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, then 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 6-11. 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 Frames and Timing MAC Headers.

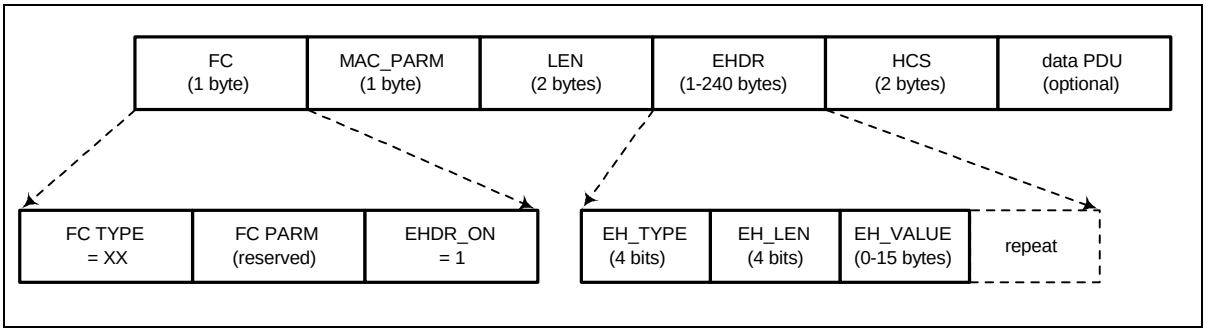


Figure 6-12. Extended MAC Format

Table 6-11. Example 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 6-12.

Table 6-12. 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 6-13 MUST be supported. Reserved and extended types are undefined at this point and MUST be ignored.

The first ten EH element types are intended for one-way transfer between the cable modem and the CMTS. The next five EH element types are for end-to-end usage within a MAC-sublayer domain. Thus, the information attached to EHDR elements 10-14 on the upstream MUST also be attached when the information is forwarded within a MAC-sublayer domain. The final EH element type is an escape mechanism that allows for more types and longer values, and MUST be as shown in Table 6-13.

Table 6-13. Extended Header Types

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]
	5	Upstream Privacy with Fragmentation[1] EH Element [DOCSIS8](See Section 6.2.7)
4 (= BP_DOWN)	4	Downstream Privacy EH Element [DOCSIS8]
5	1	Service Flow EH Element; Payload Header Suppression Header Downstream.
6	1	Service Flow EH Element; Payload Header Suppression Header Upstream
	2	Service Flow EH Element; Payload Header Suppression Header Upstream (1 byte), Unsolicited Grant Synchronization Header (1 byte)
7 - 9		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)

Notes to Table 6-13:

1. An Upstream Privacy with Fragmentation EH Element MUST only occur within a Fragmentation MAC-Specific Header. (Refer to Section 6.2.5.4)

6.2.6.1 Piggyback Requests

Several Extended Headers can be used to request bandwidth for subsequent transmissions. These requests are generically referred to as “piggyback requests”. They are extremely valuable for performance because they are not subject to contention as Request Frames generally are. (Refer to Section 7.4)

Requests for additional bandwidth can be included in Request, Upstream Privacy and Upstream Privacy with Fragmentation Extended Header elements.

6.2.6.2 Fragmentation Extended Header

Fragmented packets use a combination of the Fragmentation MAC header and a modified version of the Upstream Privacy Extended header. Section 6.2.5.4 describes the Fragmentation MAC header. The Upstream

Privacy Extended Header with Fragmentation, also known as the Fragmentation Extended Header, MUST be as shown in Table 6-14.

Table 6-14. Fragmentation Extended Header Format

EH Element Fields	Usage	Size
EH_TYPE	Upstream Privacy EH element = 3	4 bits
EH_LEN	Length of EH_VALUE = 5	4 bits
EH_VALUE	Key_seq; same as in BP_UP	4 bits
	Ver = 1; version number for this EHDR	4 bits
	BPI_ENABLE If BPI_ENABLE=0, BPI disabled If BPI_ENABLE=1, BPI enabled	1 bit
	Toggle bit; same as in BP_UP (refer to [DOCSIS8])	1 bit
	SID; Service ID associated with this fragment	14 bits
	REQ; number of mini-slots for a piggyback request	8 bits
	Reserved; must be set to zero	2 bits
	First_Frag; set to one for first fragment only	1 bit
	Last_Frag; set to one for last fragment only	1 bit
	Frag_seq; fragment sequence count, incremented for each fragment.	4 bits

6.2.6.3 Service Flow Extended Header

The Service Flow EH Element is used to enhance Service Flow operations. It may consist of one or two bytes in the EH_VALUE field. The Payload Header Suppression Header is the only byte in a one byte field or the first byte in a two byte field. The Unsolicited Grant Synchronization Header is the second byte in a two byte field.

6.2.6.3.1 Payload Header Suppression Header

In Payload Header Suppression (PHS), a repetitive portion of the payload headers following the HCS is suppressed by the sending entity and restored by the receiving entity. In the upstream, the sending entity is the

CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM.

For small payloads, Payload Header Suppression provides increased bandwidth efficiency without having to use compression. Payload Header Suppression may be separately provisioned in the upstream and downstream, and is referenced with an extended header element.

A compliant CM MUST support Payload Header Suppression.[4] A compliant CMTS MAY support Payload Header Suppression.

The Payload Header Suppression Extended Header sub-element has the following format:

Table 6-15. Payload Header Suppression EHDR Sub-Element Format

EH Element Fields	Usage		Size
EH_TYPE	Service Flow EH_TYPE=5 for downstream and EH_TYPE=6 for upstream		4 bits
EH_LEN	Length of EH_VALUE = 1		4 bits
EH_VALUE	0	Indicates no payload header suppression on current packet.	8 bits
	1-255	Payload Header Suppression Index (PHSI)	

The Payload Header Suppression Index is unique per SID in the upstream and unique per CM in the downstream. Payload Header Suppression is disabled if this Extended Header element is omitted or, if included, with the PHSI value set to 0. The Payload Header Suppression Index (PHSI) references the suppressed byte string known as a Payload Header Suppression Field (PHSF).

Note: While PHS signaling allows for up to 255 Payload Header Suppression Rules per Service Flow, the exact number of PHS rules supported per Service Flow is implementation dependent. Similarly, PHS signaling allows for PHS Sizes of up to 255 bytes, however, the maximum PHS Size supported is implementation dependent. For interoperability, the minimum PHS Size that MUST be supported is 64 bytes for any PHS rule supported. As with any other parameter requested in a Dynamic Service Request, a PHS-related DSx request can be denied because of a lack of resources.

4. This is not intended to imply that the CM must be capable of determining when to invoke Payload Header Suppression. Payload Header Suppression support is only required for the explicitly signalled case.

The Upstream Suppression Field MUST begin with the first byte following the MAC Header Checksum. The Downstream Suppression Field MUST begin with the thirteenth byte following the MAC Header Checksum. This allows the Ethernet SA and DA to be available for filtering by the CM.

The operation of Baseline Privacy (refer to [DOCSIS8]) is not affected by the use of PHS. When Fragmentation is inactive, Baseline Privacy begins encryption and decryption with the thirteenth byte following the MAC Header checksum.

Unless the entire Packet PDU is suppressed, the Packet PDU CRC is always transmitted, and MUST be calculated only on the bytes transmitted. The bytes that are suppressed MUST NOT be included in the CRC calculation.

6.2.6.3.2 Unsolicited Grant Synchronization Header

The Unsolicited Grant Synchronization Header may be used to pass status information regarding Service Flow scheduling between the CM and CMTS. It is currently only defined for use in the upstream with Unsolicited Grant and Unsolicited Grant with Activity Detection scheduling services. (Refer to Section 8.2).

This extended header is similar to the Payload Suppression EHDR except that the EH_LEN is 2, and the EH_VALUE has one additional byte which includes information related to Unsolicited Grant Synchronization. For all other Service Flow Scheduling Types, the field SHOULD NOT be included in the Extended Header Element generated by the CM. The CMTS MAY ignore this field.

Table 6-16. Unsolicited Grant Synchronization EHDR Sub-Element Format

EH Element Fields	Usage		Size
EH_TYPE	Service Flow EH_TYPE = 6		4 bits
EH_LEN	Length of EH_VALUE = 2		4 bits
EH_VALUE	0	Indicates no payload header suppression on current packet.	8 bits [always present]
	1-255	Payload Header Suppression Index (PHSI)	
	Queue Indicator		1 bit
	Active Grants		7 bits

6.2.7 Fragmented MAC Frames

When enabled, fragmentation is initiated any time the grant length is less than the requested length. This normally occurs because the CMTS chooses to grant less than the requested bandwidth.

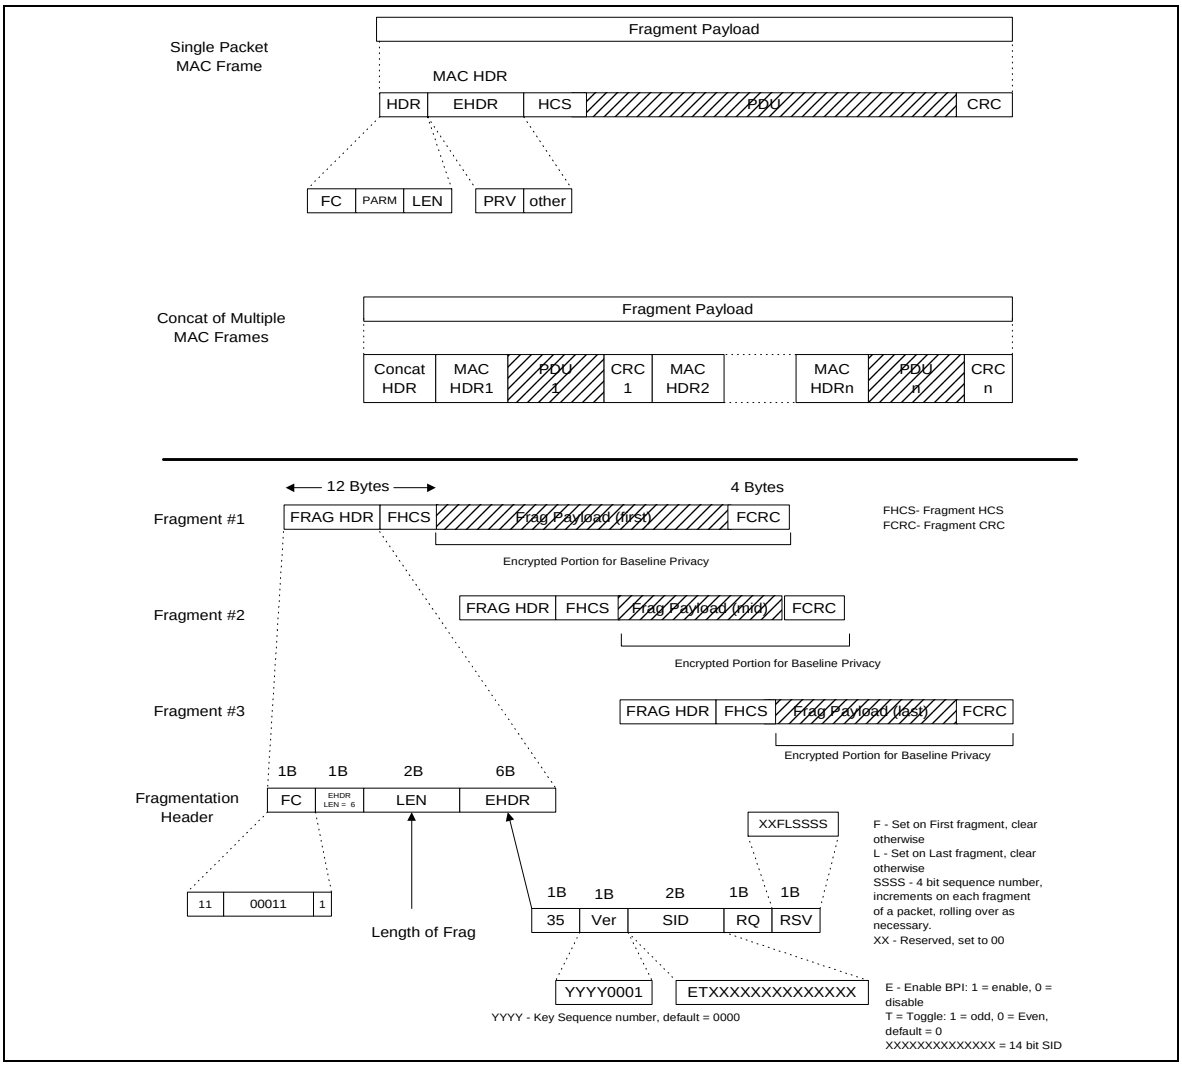


Figure 6-13. Fragmentation Details

The CM MAC calculates how many bytes of the original frame, including overhead for a fragmentation header and CRC, can be sent in the received grant. The CM MAC generates a fragmentation header for each fragment. Fragmented frames use the MAC Message type (FC = 11). The FC parameter field is set to (00011), in order to uniquely identify the fragmentation header from other MAC Message types. A four bit sequence field is used in the last byte of the Extended Header field to aid in reassembly and to detect dropped or missing fragments. The CM arbitrarily selects a sequence number for the first fragment of a frame.[5] Once the sequence number is selected for the first fragment, the CM MUST increment the sequence number by one for each fragment transmitted for that frame. There are two flags associated with the sequence number, F and L, where F is set to indicate the first fragment and L is set to indicate the last fragment. Both are cleared for middle fragments. The

5. Note, 'frame' always refers to either frames with a single Packet PDU or concatenated frame.

CMTS stores the sequence number of the first fragment (F bit set) of each frame. The CMTS MUST verify that the fragment sequence field increments (by one) for each fragment of the frame.

The REQ field in the fragmentation header is used by the fragmentation protocol for First and Middle fragments (refer to Section 8.3). For the Last fragment, the REQ field is interpreted as a request for bandwidth for a subsequent frame.

Fragmentation headers are fixed size and MUST contain only a Fragmentation extended header element. The extended header consists of a Privacy EH element extended by one byte to make the fragment overhead an even 16 bytes. A Privacy EH element is used whether the original packet header contained a Privacy EH element or not. If privacy is in use, the following fields, Version, Enable bit, and SID, in the fragment EH element are the same with those of BP EH element inside the original MAC frame. If privacy is not in use, the Privacy EH element is used but the enable bit is cleared. The SID used in the fragment EH element MUST match the SID used in the Partial Grant that initiated the fragmentation. A separate CRC MUST be calculated for each fragment (note that each MAC frame payload will also contain the CRC for that packet). A packet CRC of a reassembled packet MAY be checked by the CMTS even though an FCRC covers each fragment.

The CMTS MUST make certain that any fragmentary grant it makes is large enough to hold at least 17 bytes of MAC layer data. This is to ensure that the grant is large enough to accommodate fragmentation overhead plus at least 1 byte of actual data. The CMTS may want to enforce an even higher limit as small fragments are extremely inefficient.

When Fragmentation is active, Baseline Privacy encryption and decryption begin with the first byte following the MAC Header checksum.

6.2.7.1 Considerations for Concatenated Packets and Fragmentation

MAC Management Messages and Data PDUs can occur in the same concatenated frame. Without fragmentation, the MAC Management Messages within a concatenated frame would be unencrypted. However, with fragmentation enabled on the concatenated frame, the entire concatenated frame is encrypted based on the Privacy Extended Header Element. This allows Baseline Privacy to encrypt each fragment without examining its contents. Clearly, this only applies when Baseline Privacy is enabled.

To ensure encryption synchronization, if fragmentation, concatenation and Baseline Privacy are all enabled, a CM MUST NOT concatenate BPKM MAC Management messages. This ensures that BPKM MAC Management messages are always sent unencrypted.

6.2.8 Error-Handling

The cable network is a potentially harsh environment that can cause several different error conditions to occur. This section, together with Section 9.5, 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 can 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.

6.2.8.1 Error Recovery During Fragmentation

There are some special error handling considerations for fragmentation. Each fragment has its own fragmentation header complete with an HCS and its own FCRC. There may be other MAC headers and CRCs within the fragmented payload. However, only the HCS of the fragment header and the FCRC are used for error detection during fragment reassembly.

If the HCS for a fragment fails the CMTS MUST discard that fragment. If the HCS passes but the FCRC fails, the CMTS MUST discard that fragment, but MAY process any requests in the fragment header. The CMTS SHOULD process such a request if it is performing fragmentation in Piggyback Mode. (Refer to Section 8.3.2.2) This allows the remainder of the frame to be transmitted as quickly as possible.

If a CMTS is performing fragmentation in Multiple Grant Mode (refer to Section 8.3.2.1) it SHOULD complete all the grants necessary to fulfil the CM's original request even if a fragment is lost or discarded. This allows the remainder of the frame to be transmitted as quickly as possible.

If any fragment of a non-concatenated MAC frame is lost or discarded the CMTS MUST discard the rest of that frame. If a fragment of a concatenated MAC frame is lost or discarded the CMTS MAY forward any frames within the concatenation that have been received correctly or it MAY discard all the frames in the concatenation.

A CMTS MUST terminate fragment reassembly if any of the following occurs for any fragment on a given SID:

- The CMTS receives a fragment with the L bit set.
- The CMTS receives an upstream fragment, other than the first one, with the F bit set.
- The CMTS receives a packet PDU frame with no fragmentation header.
- The CMTS deletes the SID for any reason.

In addition, the CMTS MAY terminate fragment reassembly based on implementation dependent criteria such as a reassembly timer. When a CMTS terminates fragment reassembly it MUST dispose of (either by discarding or forwarding) the reassembled frame(s).

6.2.8.2 Error Codes and Messages

SP-OSSIv1.1 [DOCSIS5] Appendix F lists CM and CMTS error codes and messages. When reporting error conditions, 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 6-14. Figure 6-14 shows the MAC Header and the MAC Management Message Header fields which are common across all MAC Management Messages.

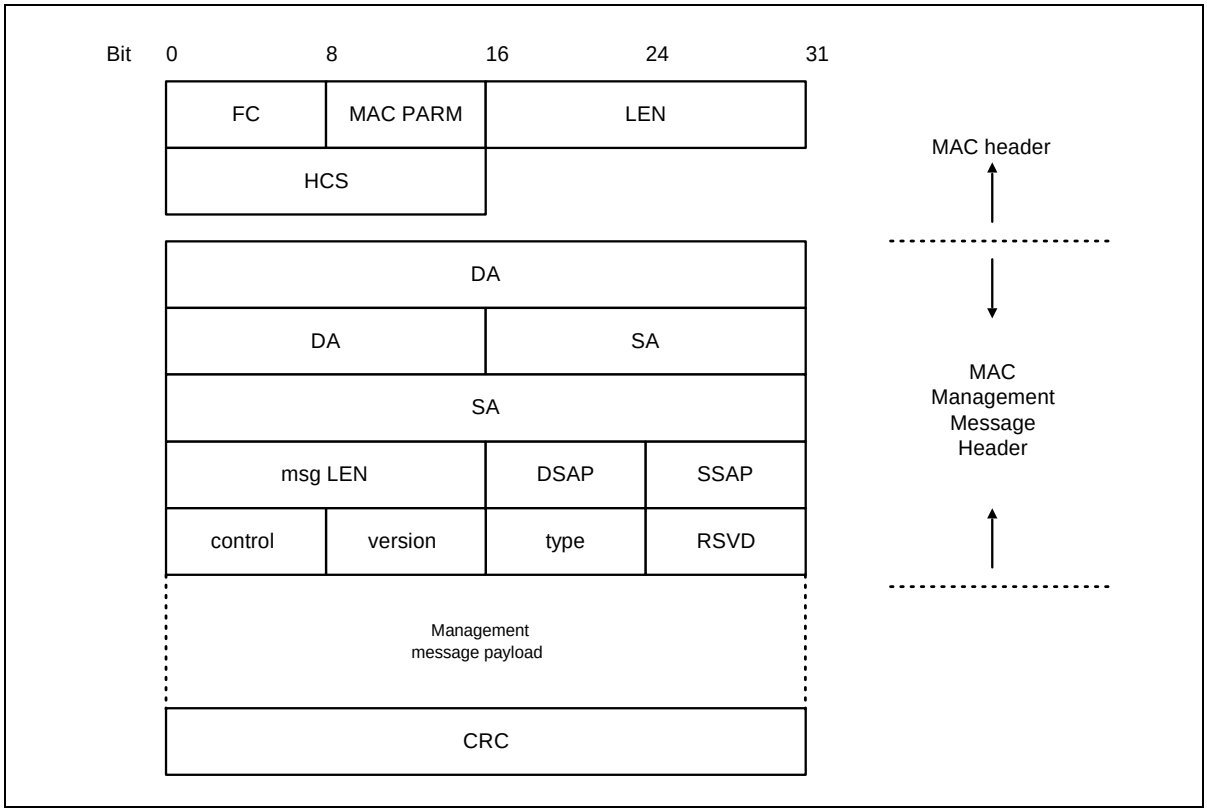


Figure 6-14. MAC Header and MAC Management Message Header Fields

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 destination SAP (00) as defined by [ISO8802-2].

SSAP The LLC null source SAP (00) as defined by [ISO8802-2].

Control Unnumbered information frame (03) as defined by [ISO8802-2].

Version & Type Each 1 octet. Refer to Table 6-17.

Table 6-17. MAC Management Message Types

Type Value	Version	Message Name	Message Description
1	1	SYNC	Timing Synchronization
2	1	UCD	Upstream Channel Descriptor
3	1	MAP	Upstream Bandwidth Allocation
4	1	RNG-REQ	Ranging Request
5	1	RNG-RSP	Ranging Response
6	1	REG-REQ	Registration Request
7	1	REG-RSP	Registration Response
8	1	UCC-REQ	Upstream Channel Change Request
9	1	UCC-RSP	Upstream Channel Change Response
10	1	TRI-TCD	Telephony Channel Descriptor [DOCSIS6]
11	1	TRI-TSI	Termination System Information [DOCSIS6]
12	1	BPKM-REQ	Privacy Key Management Request [DOCSIS8]]
13	1	BPKM-RSP	Privacy Key Management Response [DOCSIS8]
14	2	REG-ACK	Registration Acknowledge
15	2	DSA-REQ	Dynamic Service Addition Request
16	2	DSA-RSP	Dynamic Service Addition Response
17	2	DSA-ACK	Dynamic Service Addition Acknowledge
18	2	DSC-REQ	Dynamic Service Change Request
19	2	DSC-RSP	Dynamic Service Change Response
20	2	DSC-ACK	Dynamic Service Change Acknowledge
21	2	DSD-REQ	Dynamic Service Deletion Request
22	2	DSD-RSP	Dynamic Service Deletion Response
23	2	DCC-REQ	Dynamic Channel Change Request
24	2	DCC-RSP	Dynamic Channel Change Response
25	2	DCC-ACK	Dynamic Channel Change Acknowledge
26	2	DCI-REQ	Device Class Identification Request
27	2	DCI-RSP	Device Class Identification Response
28	2	UP-DIS	Upstream Transmitter Disable
29–255			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].

A compliant CMTS or CM MUST support the MAC management message types listed in Table 6-17, except messages specific to Telephony Return devices which MAY be supported.

6.3.2 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 with FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header. This MUST be followed by a Packet PDU in the format shown in Figure 6-15.

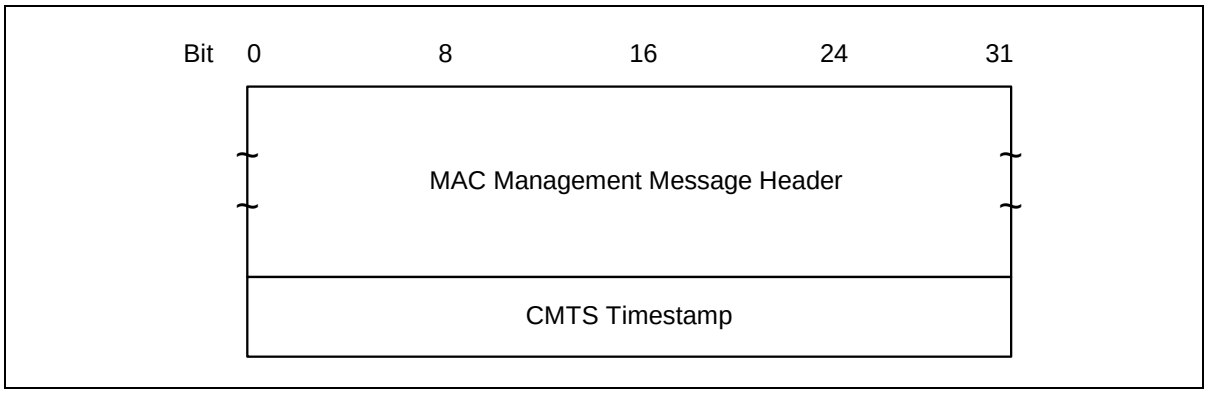


Figure 6-15. Format of Packet PDU Following the Timing Header

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 a SYNC message to cross an MPEG packet boundary.[6]

[6] 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 one half the smallest possible mini-slot at the highest possible symbol rate. See Section 7.3.4 for time-unit relationships.

6.3.3 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 6-16). 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.

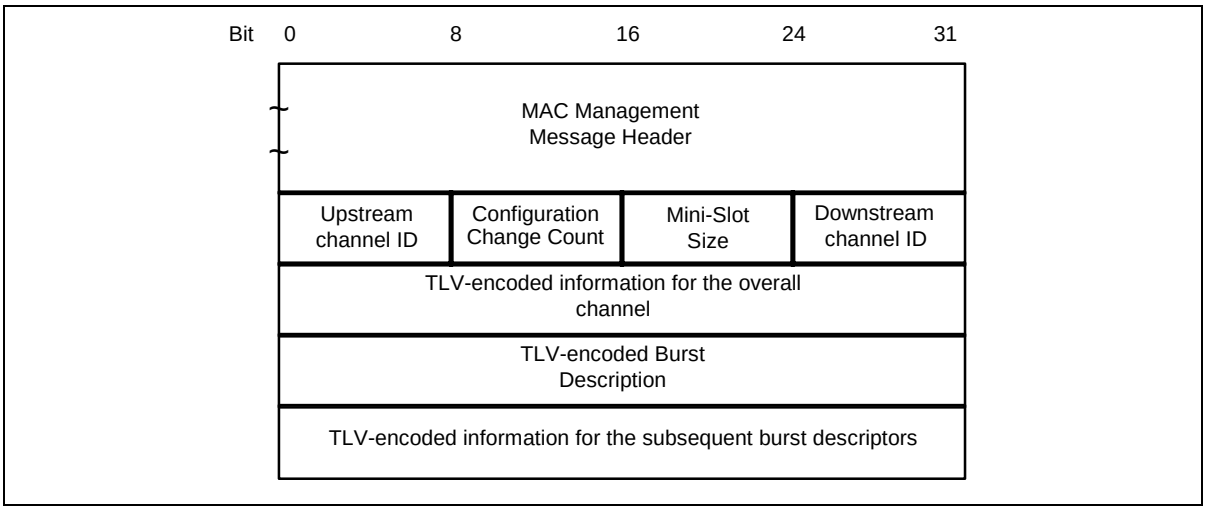


Figure 6-16. Upstream Channel Descriptor

A CMTS MUST generate UCDs in the format shown in Figure 6-16, 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 µs. Allowable values are $T = 2^M$, M = 1,...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. Note: Upstream Channel ID = 0 is reserved to indicate telephony return [DOCSIS6].

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 Table 6-18, for channel parameters, and Table 6-19, for upstream physical layer burst attributes. Channel-wide parameters (types 1-3 in Table 6-18) MUST preceed burst descriptors (type 4 below).

Table 6-18. 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 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 6.3.4 and Table 6-20). The format of the Burst Descriptor is shown in Figure 6-17.

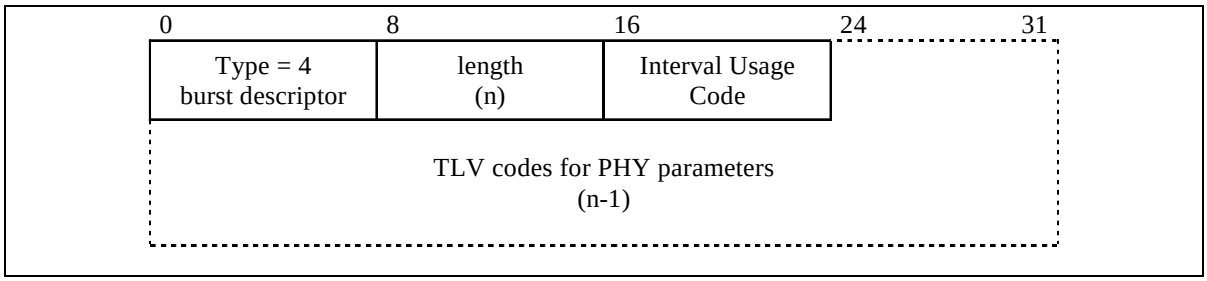


Figure 6-17. Top-Level Encoding for a Burst Descriptor

In Figure 6-17:

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 6-20. 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 6-19.

A Burst Descriptor MUST be included for each Interval Usage Code that is to be used in the allocation MAP. The Interval Usage Code MUST be one of the values from Table 6-20.

Within each Burst Descriptor is an unordered list of Physical-layer attributes, encoded as TLV values. These attributes are shown in Table 6-19.

Table 6-19. 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 6-18). 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 4-9), and is I_1 in the first symbol of the burst (see Section 4.2.2.2).
FEC Error Correction (T)	5	1	0-10 (0 implies no FEC. The number of codeword parity bytes is 2*T)
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 Section 7.1.2.5) If the CMTS needs to limit the maximum length of concatenated frames it SHOULD use this configuration setting to do so.

Guard Time Size	9	1	Number of symbol times measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. In all burst descriptors, the CMTS MUST choose the parameters such that the number of bytes that fit into any valid number of mini-slots will not change if the guard time is decreased by 1.
Last Codeword Length	10	1	1 = fixed; 2 = shortened
Scrambler on/off	11	1	1 = on; 2 = off

6.3.3.1 Example of UCD Encoded TLV Data

An example of UCD encoded TLV data is given in Figure 6-18.

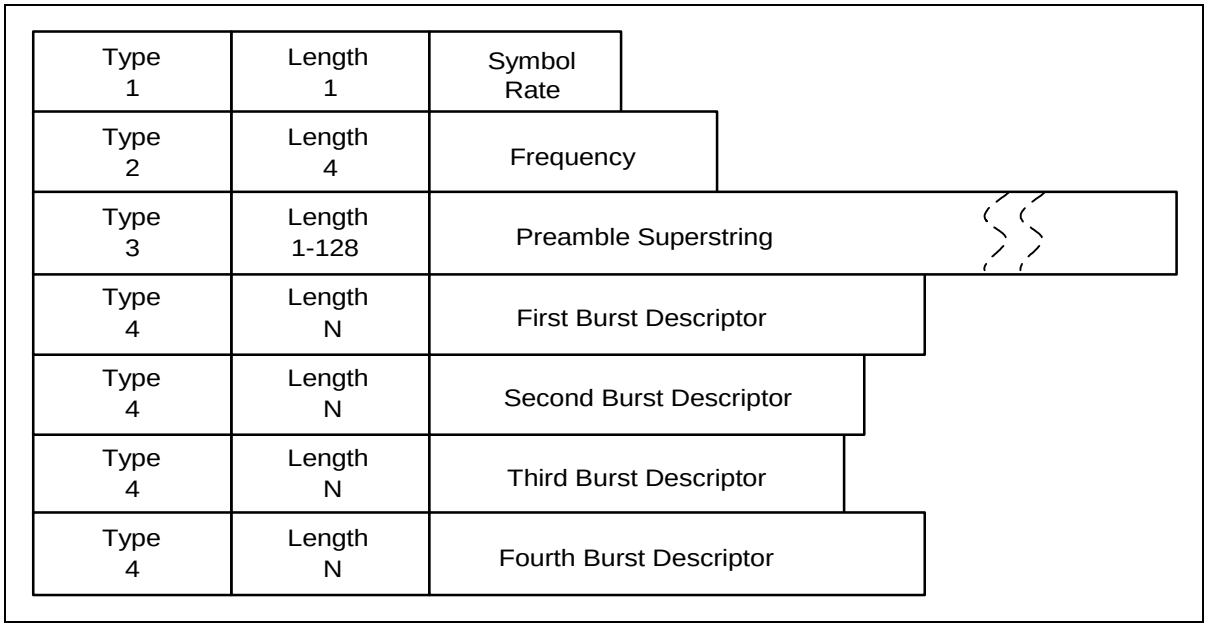


Figure 6-18. Example of UCD Encoded TLV Data

6.3.4 Upstream Bandwidth Allocation Map (MAP)

A CMTS MUST generate MAPs in the format shown in Figure 6-19.

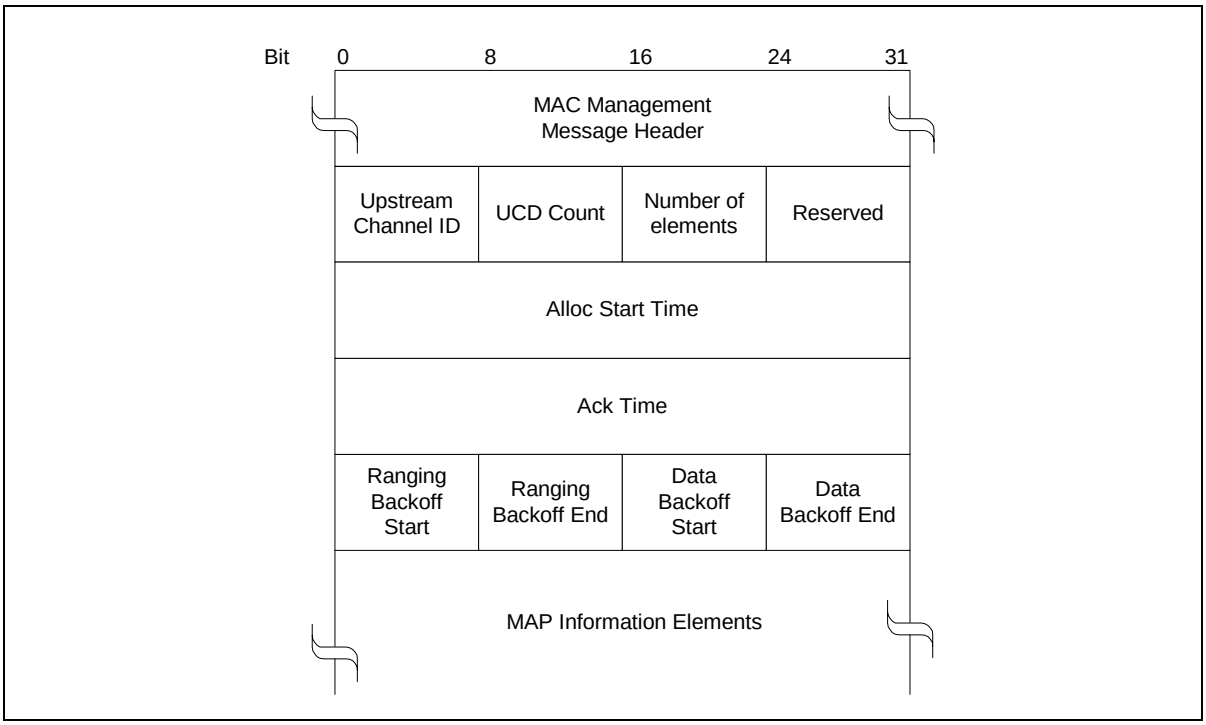


Figure 6-19. MAP Format

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 9.3.2.
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 7.4.
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 6-20 and Table 6-20. Values for IUCs are defined in Table 6-20 and are described in detail in Section 7.1.2.

Note: The lower (26-M) bits of the Alloc Start Time and Ack Time MUST be used as the effective MAP start and ack times where M is given in Section 6.3.3. The relationship between the Alloc Start/Ack time counters and the timestamp counter is described in Section 7.4.

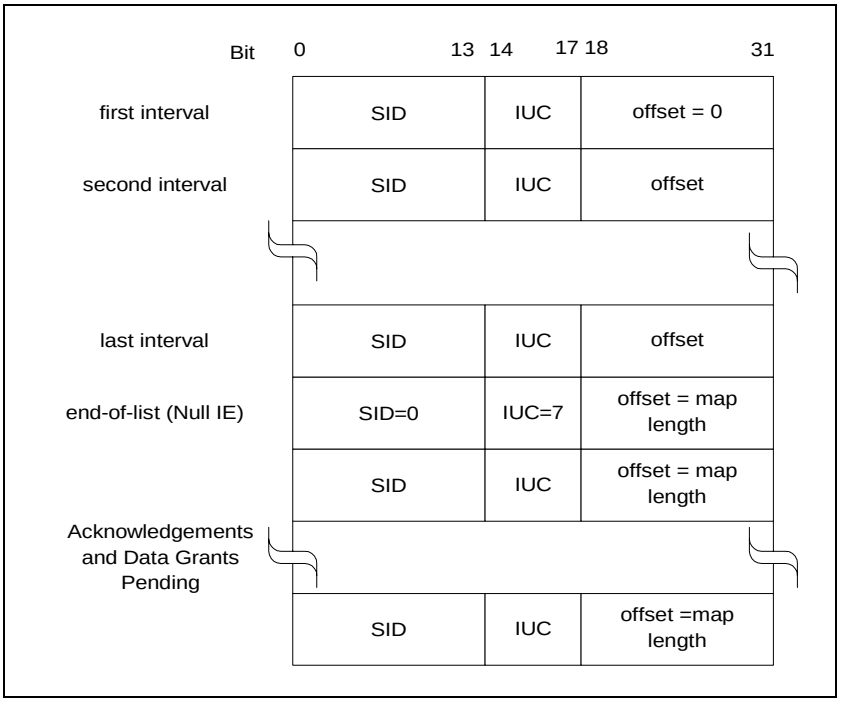


Figure 6-20. MAP Information Element Structure

Table 6-20. Allocation MAP Information Elements (IE)

IE Name[1]	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	Starting offset of MAINT region (used in Initial Ranging)
Station Maintenance[2]	4	unicast[3]	Starting offset of MAINT region (used in Periodic Ranging)
Short Data Grant[4]	5	unicast	Starting offset of Data Grant assignment; If inferred length = 0, then it is a Data Grant pending.
Long Data Grant	6	unicast	Starting offset of Data Grant assignment; If inferred length = 0, then 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

Notes to Table 6-20:

1. 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.

2. 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 6-19).
3. 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.
4. 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.

6.3.5 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 and request power adjustment. This message MUST use an FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header. This MUST be followed by a Packet PDU in the format shown in Figure 6-21.

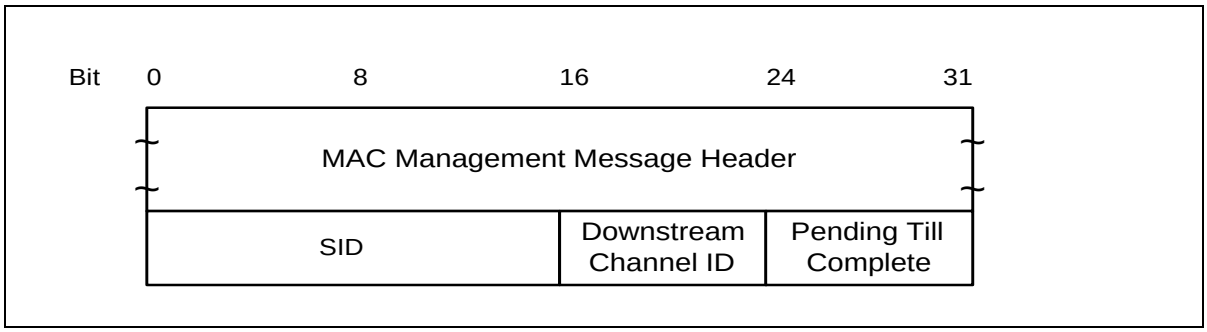


Figure 6-21. Packet PDU Following the Timing Header

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.

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, then all previous Ranging Response attributes have been applied prior to transmittting this request. If nonzero then this is 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).

6.3.6 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 9.2.4. 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.

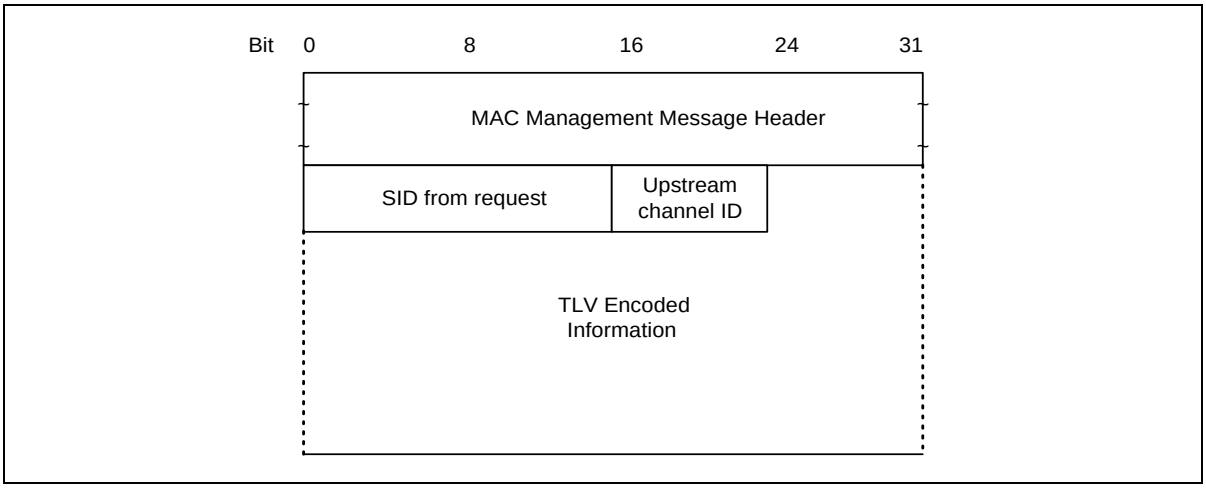


Figure 6-22. Ranging Response

A CMTS MUST generate Ranging Responses in the form shown in Figure 6-22, including all of the following parameters:

SID — If the modem is being instructed by this response to move to a different channel, this is initialization SID. Otherwise, this is the SID from the corresponding RNG-REQ to which this response refers, except that if the corresponding RNG-REQ was an initial ranging request specifying a initialization SID, then 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 H). Thus,

the CM MUST use this channel ID for the rest of its transactions, not the channel ID it initiated the range request from.

All other parameters are coded as TLV tuples.

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

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 This provides the equalization coefficients for the pre-equalizer.

Downstream Frequency Override An optional parameter. The downstream frequency with which the modem should redo initial ranging. (See Section 6.3.6.3)

Upstream Channel ID Override An optional parameter. The identifier of the upstream channel with which the modem should redo initial ranging. (See Section 6.3.6.3)

6.3.6.1 Encodings

The type values used MUST be those defined in Table 6-21 and Figure 6-23. 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.

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

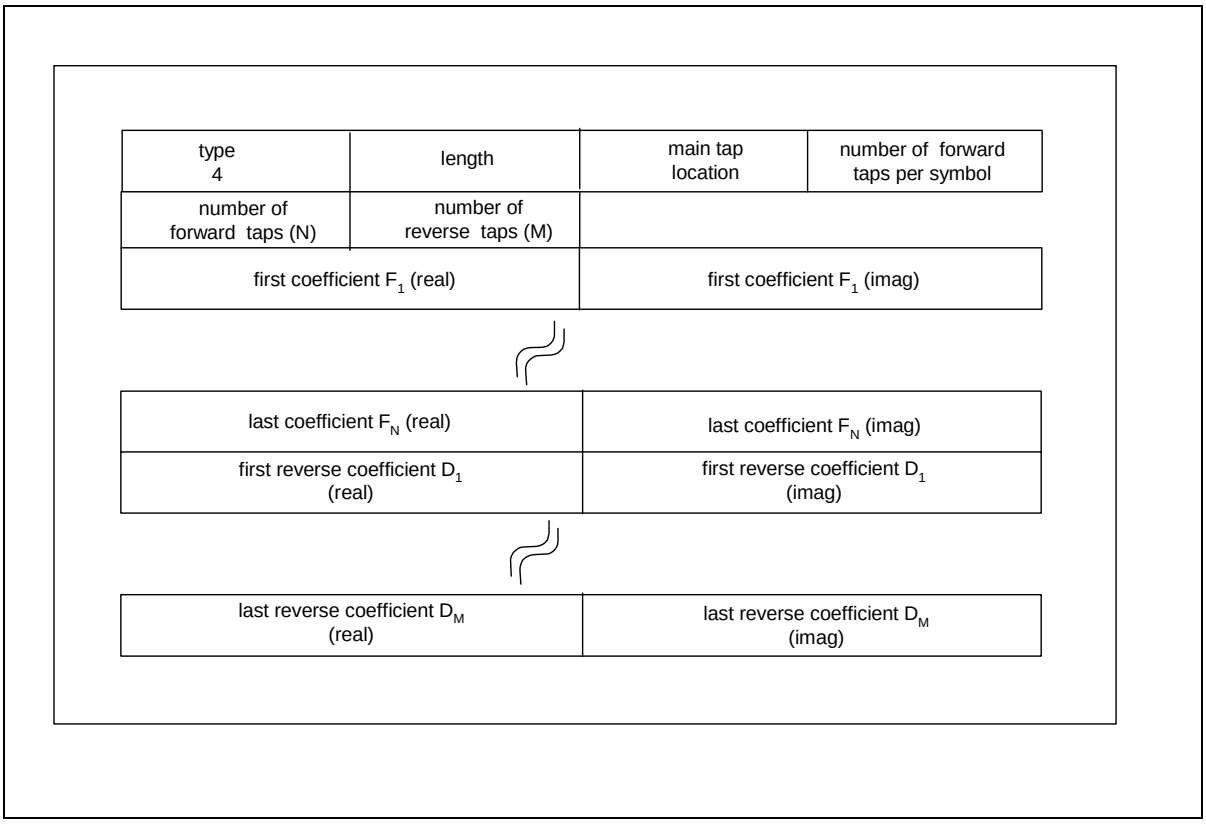


Figure 6-23. Generalized Decision Feedback Equalization Coefficients

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.

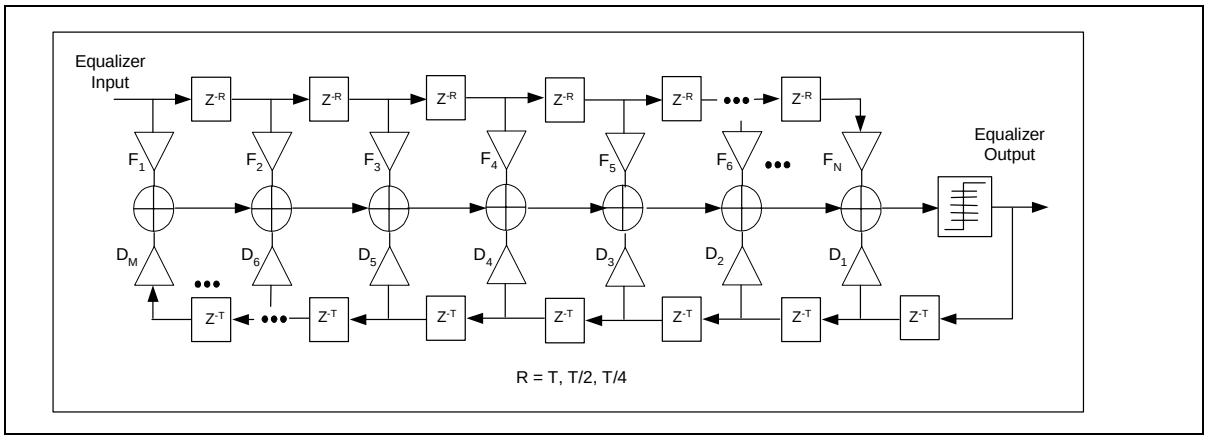


Figure 6-24. Generalized Equalizer Tap Location Definition

6.3.6.2 Example of TLV Data

An example of TLV data is given in Figure 6-25.

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 6-25. Example of TLV Data

6.3.6.3 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. If a downstream frequency override is specified in the RNG-RSP, the modem MUST reinitialize its MAC (Section 9.2) 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 9.2) 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 9.2) 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 for upstream channel ID and downstream frequency are optional, but if specified in the config file they take precedence over the ranging response parameters. Once ranging is complete, only the Appendix C.1.1.2, UCC-REQ, and DCC-REQ mechanisms are available for moving the modem to a new upstream channel, and only the Appendix C.1.1.1 and DCC-REQ mechanisms are available for moving the modem to a new downstream channel.

6.3.7 Registration Request (REG-REQ)

A Registration Request MUST be transmitted by a CM at initialization after receipt of a CM parameter file, except as outlined in Section 9.2.8 and 9.2.9.

To provide for flexibility, the message parameters following the SID MUST be encoded in a type/length/value form.

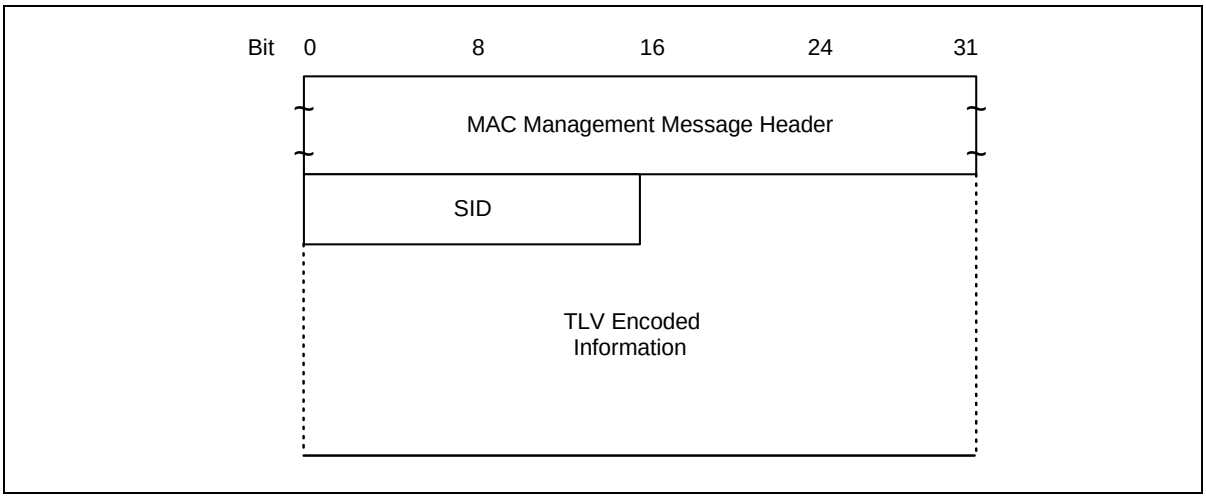


Figure 6-26. Registration Request

A CM MUST generate Registration Requests in the form shown in Figure 6-26, 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:

- All configuration settings included in the CMTS MIC calculation as specified in Appendix D.3.1
- CMTS MIC Configuration Setting

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.

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 parameter MUST also be included in the Registration Request.

- Modem Capabilities Encodings[7]

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

- 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
- SNMPv3 Kickstart Value
- CPE Ethernet MAC Address
- HMAC Digest
- End Configuration Setting
- Pad Configuration Setting
- Telephone Settings Option
- SNMPv3 Notification Receiver

7. The CM MUST specify all of its Modem Capabilities in its Registration Request subject to the restrictions in section C.1.3.1. The CMTS MUST NOT assume any Modem Capability which is defined but not explicitly indicated in the CM's Registration Request.

6.3.8 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.

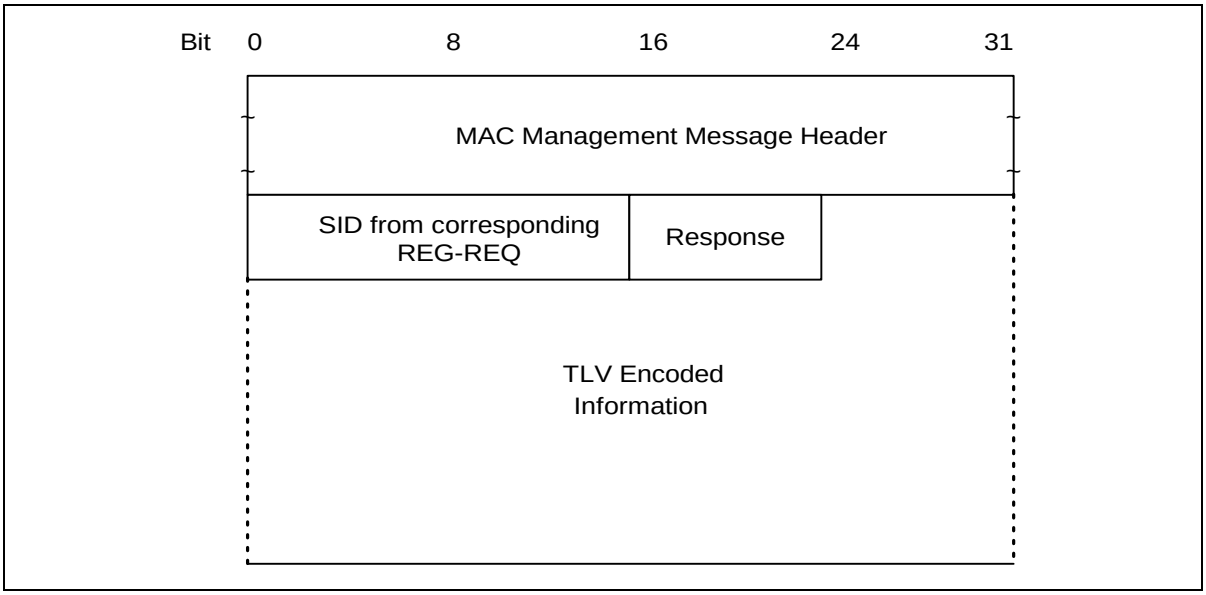


Figure 6-27. Registration Response Format

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. (This acts as a transaction identifier)

Response — For REG-RSP to a modem registering as a 1.0 modem (*i.e.*, REG-REQ contains DOCSIS 1.0 Class of Service Encodings)
0 = Okay
1 = Authentication Failure
2 = Class of Service Failure
For REG-RSP to a modem registering as a 1.1 modem (*i.e.*, REG-REQ contains Service Flow Encodings), this field MUST contain one of the Confirmation Codes in Appendix C and C.4.2.

Note: Failures apply to the entire Registration Request. Even if only a single requested Service Flow or DOCSIS 1.0 Service Class is invalid or undeliverable the entire registration is failed.

If the REG-REQ was successful, and contained Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, the REG-RSP MUST contain, for each of these:

Classifier Parameters — All of the Classifier Parameters from the corresponding REG-REQ, plus the Classifier Identifier assigned by the CMTS.

Service Flow Parameters All the Service Flow Parameters from the REG-REQ, plus the Service Flow ID assigned by the CMTS. Every Service Flow that contained a Service Class Name that was admitted/activated MUST be expanded into the full set of TLVs defining the Service Flow. Every upstream Service Flow that was admitted/activated MUST have a Service Identifier assigned by the CMTS. A Service Flow that was only provisioned will include only those QoS parameters that appeared in the REG-REQ, plus the assigned Service Flow ID.

Payload Header Suppression Parameters All the Payload Header Suppression Parameters from the REG-REQ, plus the Payload Header Suppression Index assigned by the CMTS.

If the REG-REQ failed due to Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, and the Response is not one of the major error codes in Appendix C.4.2, the REG-RSP MUST contain at least one of the following:

Classifier Error Set A Classifier Error Set and identifying Classifier Reference and Service Flow Reference MUST be included for at least one failed Classifier in the corresponding REG-REQ. Every Classifier Error Set MUST include at least one specific failed Classifier Parameter of the corresponding Classifier.

Service Flow Error Set A Service Flow Error Set and identifying Service Flow Reference MUST be included for at least one failed Service Flow in the corresponding REG-REQ. Every Service Flow Error Set MUST include at least one specific failed QoS Parameter of the corresponding Service Flow.

Payload Header Suppression Error Set A PHS Error Set and identifying Service Flow Reference and Classifier Reference pair MUST be included for at least one failed PHS Rule in the corresponding REG-REQ. Every PHS Error Set MUST include at least one specific failed PHS Parameter of the corresponding failed PHS Rule.

Service Class Name expansion always occurs at admission time. Thus, if a Registration-Request contains a Service Flow Reference and a Service Class Name for deferred admission/activation, the Registration-Response MUST NOT include any additional QoS Parameters except the Service Flow Identifier. (Refer to Section 8.1.3)

If the corresponding Registration Request contains DOCSIS 1.0 Service Class TLV's (refer to Appendix C.1.1.4), the Registration Response MUST contain the following TLV tuples:

DOCSIS 1.0 Service Class Data Returned when Response = Okay
Service ID / service class tuple for each class of service granted

Note: Service class IDs MUST be those requested in the corresponding REG-REQ.

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.

All other parameters are coded TLV tuples

Modem Capabilities The CMTS response to the capabilities of the modem (if present in the Registration Request)

Vendor-Specific Data As defined in Appendix C

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

6.3.8.1 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.

6.3.8.1.1 Modem Capabilities

This field defines the CMTS response to the modem capability field in the Registration Request. The CMTS MUST respond to the modem capability to indicate whether they may be used. If the CMTS does not recognize a modem capability, it MUST return the TLV with the value zero ("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. Capabilities set to "off" in the REG-REQ MUST also be set to "off" in the REG-RSP.

Encodings are as defined for the Registration Request.

6.3.8.1.2 DOCSIS 1.0 Service Class Data

A DOCSIS 1.0 Service Class Data parameter MUST be present in the Registration Response for each DOCSIS 1.0 Class of Service parameter (refer to Appendix C.1.1.4) in the Registration Request.

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.

Each received DOCSIS 1.0 Class of Service parameter must have a unique Class ID in the range 1..16. If no Class ID was present for any single DOCSIS 1.0 Class-of-Service TLV in the REG-REQ, the CMTS MUST send a REG-RSP with a class-of-service failure response and no DOCSIS 1.0 Class-of-Service TLVs.

Type	Length	Value
1	n	Encoded service class data

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.

Type	Length	Value
1.1	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
1.2	2	SID

6.3.9 Registration Acknowledge (REG-ACK)

A Registration Acknowledge MUST be transmitted by the CM in response to a REG-RSP from the CMTS with a confirmation code of ok (0).[8] It confirms acceptance by the CM of the QoS parameters of the flow as reported by the CMTS in it REG-RSP. The format of a REG-ACK MUST be as shown in Figure 6-28.

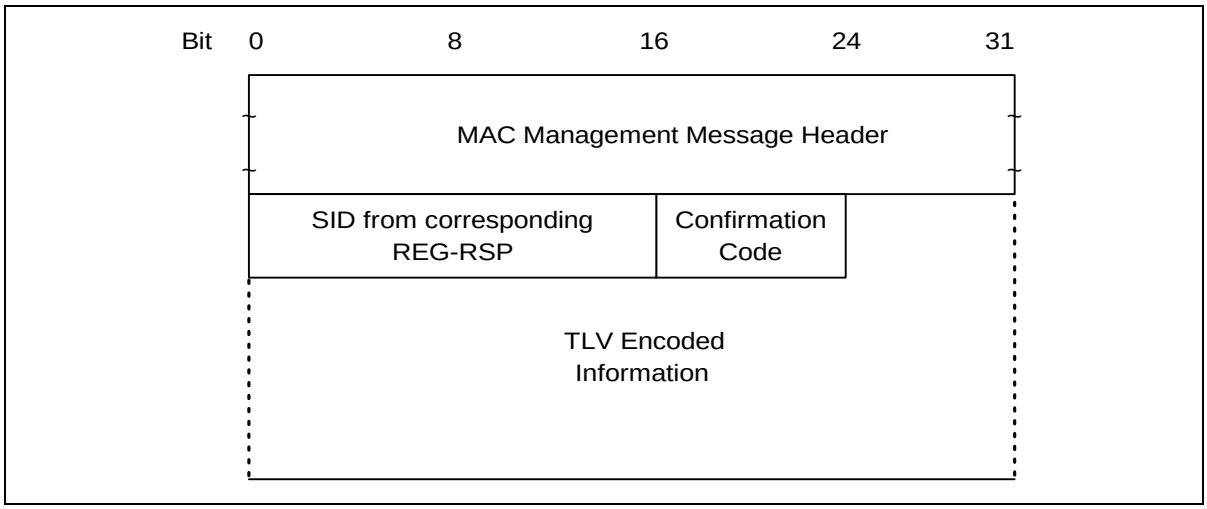


Figure 6-28. Registration Acknowledgment

The parameter MUST be as follows:

SID from Corresponding REG-RSP SID from corresponding REG-RSP to which this acknowledgment refers. (This acts as a transaction identifier)

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the entire corresponding Registration Response.

The CM is required to send all provisioned Classifiers, Service Flows and Payload Header Suppression Rules to the CMTS in the REG-REQ (see section 6.3.7). The CMTS will return them with Identifiers, expanding Service Class Names if present, in the REG-RSP (see Section 6.3.8). Since the CM may be unable to support one or more of these provisioned items, the REG-ACK includes Error Sets for all failures related to these provisioned items.

If there were any failures of provisioned items, the REG-ACK MUST include the Error Sets corresponding to those failures. The Error Set identification is provided by using Service Flow ID and Classifier ID from

8. The Registration-Acknowledge is a DOCSIS 1.1 message. Refer to Appendix G for details of registration interoperability issues.

corresponding REG-RSP. If a Classifier ID or SFID was omitted in the REG-RSP, the CM MUST use the appropriate Reference (Classifier Reference, SF Reference) in the REG-ACK.

Classifier Error Set — A Classifier Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for at least one failed Classifier in the corresponding REG-RSP. Every Classifier Error Set MUST include at least one specific failed Classifier Parameter of the corresponding Classifier. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.

Service Flow Error Set — A Service Flow Error Set of the REG-ACK message encodes specifics of failed Service Flows in the REG-RSP message. A Service Flow Error Set and identifying Service Flow Reference/Identifier MUST be included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding REG-RSP message. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.

Payload Header Suppression Error Set — A PHS Error Set and identifying Service flow Reference/Identifier and Classifier Reference/Identifier pair MUST be included for at least one failed PHS Rule in the corresponding REG-RSP. Every PHS Error Set MUST include at least one specific failed PHS of the failed PHS Rule. This parameter MUST be omitted if the entire REG-REQ/RSP is successful.

Per Service Flow acknowledgment is necessary not just for synchronization between the CM and CMTS, but also to support use of the Service Class Name. (Refer to Section 8.1.3) Since the CM may not know all of the Service Flow parameters associated with a Service Class Name when making the Registration Request, it may be necessary for the CM to NAK a Registration Response if it has insufficient resources to actually support this Service Flow.

6.3.10 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.

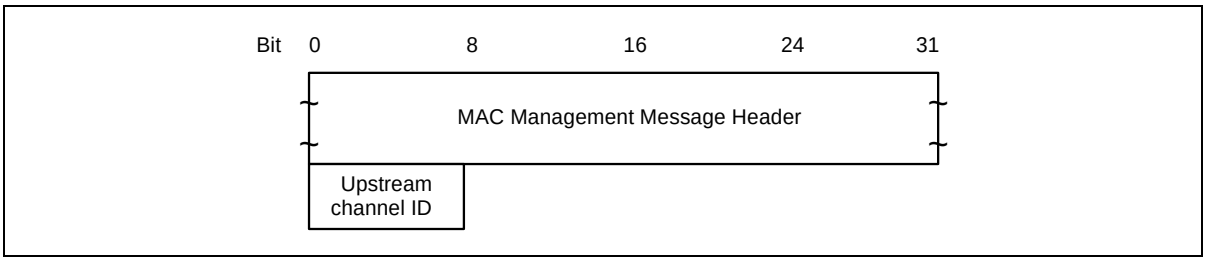


Figure 6-29. Upstream Channel Change Request

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.

6.3.11 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 a 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.

After switching to a new upstream channel, a CM MUST re-range using broadcast initial ranging, and then MUST proceed without re-performing registration. The full procedure for changing channels is described in Section 9.3.3.

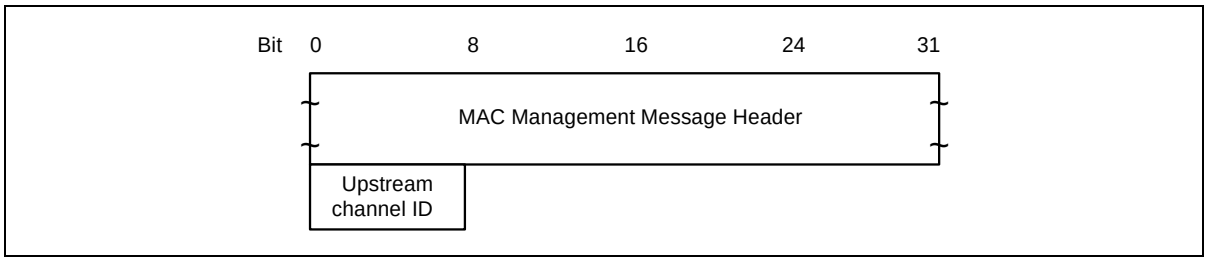


Figure 6-30. Upstream Channel Change Response

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 MUST be the same Channel ID specified in the UCC-REQ message. This MUST be an 8-bit field.

6.3.12 Dynamic Service Addition — Request (DSA-REQ)

A Dynamic Service Addition Request MAY be sent by a CM or CMTS to create a new Service Flow.

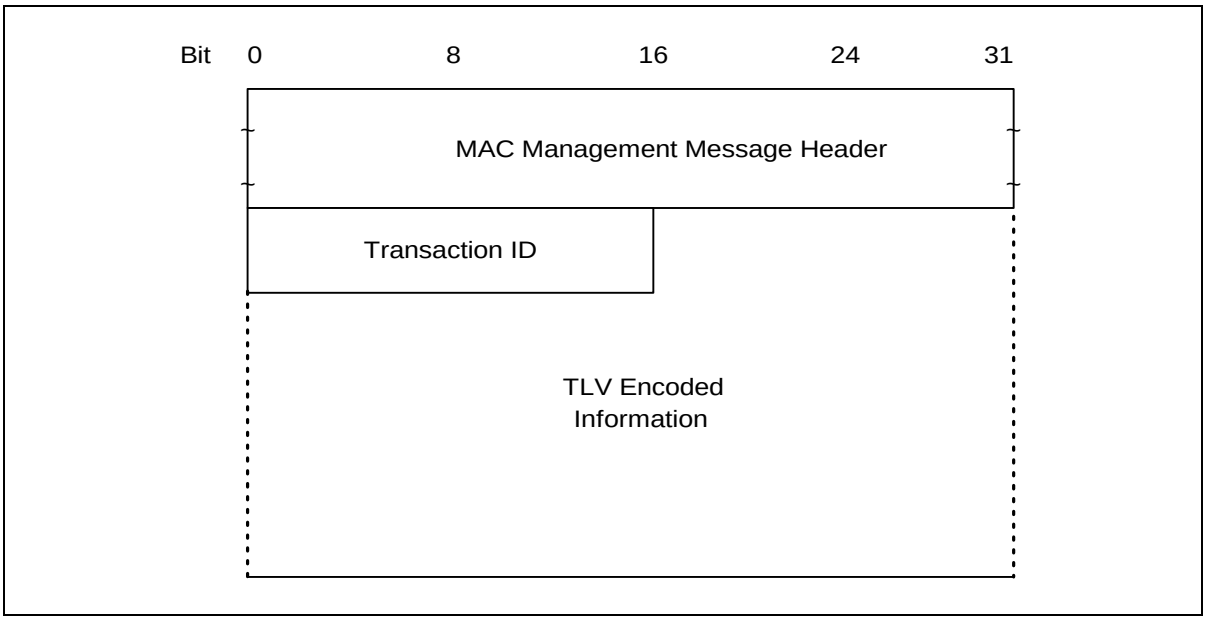


Figure 6-31. Dynamic Service Addition — Request

A CM or CMTS MUST generate DSA-REQ messages in the form shown in Figure 6-31 including the following parameter:

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C. A DSA-REQ message MUST NOT contain parameters for more than one Service Flow in each direction, *i.e.*, a DSA-REQ message MUST contain parameters for either a single upstream Service Flow, or for a single downstream Service Flow, or for one upstream and one downstream Service Flow.

The DSA-REQ message MUST contain:

Service Flow Parameters Specification of the Service Flow's traffic characteristics and scheduling requirements.

The DSA-REQ message MAY contain classifier parameters and payload header suppression parameters associated with the Service Flows specified in the message:

Classifier Parameters Specification of the rules to be used to classify packets into a specific Service Flow.

Payload Header Suppression Parameters Specification of the payload header suppression rules to be used with an associated classifier.

If Privacy is enabled, the DSA-REQ message MUST contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.12.1 CM-Initiated Dynamic Service Addition

CM-initiated DSA-Requests MUST use the Service Flow Reference to link Classifiers to Service Flows. Values of the Service Flow Reference are local to the DSA message; each Service Flow within the DSA-Request MUST be assigned a unique Service Flow Reference. This value need not be unique with respect to the other service flows known by the sender.

CM-initiated DSA-Request MUST use the Classifier Reference and Service Flow Reference to link Payload Header Suppression Parameters to Classifiers and Service Flows. A DSA-request MUST use the Service Flow

Reference to link Classifier to Service Flow. Values of the Classifier Reference are local to the DSA message; each Classifier within the DSA-request MUST be assigned a unique Classifier Reference.

CM-initiated DSA-Requests MAY use the Service Class Name (refer to Appendix C.2.2.3.4) in place of some, or all, of the QoS Parameters.

6.3.12.2 CMTS-Initiated Dynamic Service Addition

CMTS-initiated DSA-Requests MUST use the Service Flow ID to link Classifiers to Service Flows. Service Flow Identifiers are unique within the MAC domain. CMTS-initiated DSA-Requests for Upstream Service Flows MUST also include a Service ID.

CMTS-initiated DSA-Requests which include Classifiers, MUST assign a unique Classifier Identifier on a per Service Flow basis.

CMTS-initiated DSA-Requests for named Service Classes MUST include the QoS Parameter Set associated with that Service Class.

6.3.13 Dynamic Service Addition — Response (DSA-RSP)

A Dynamic Service Addition Response MUST be generated in response to a received DSA-Request. The format of a DSA-RSP MUST be as shown in Figure 6-32.

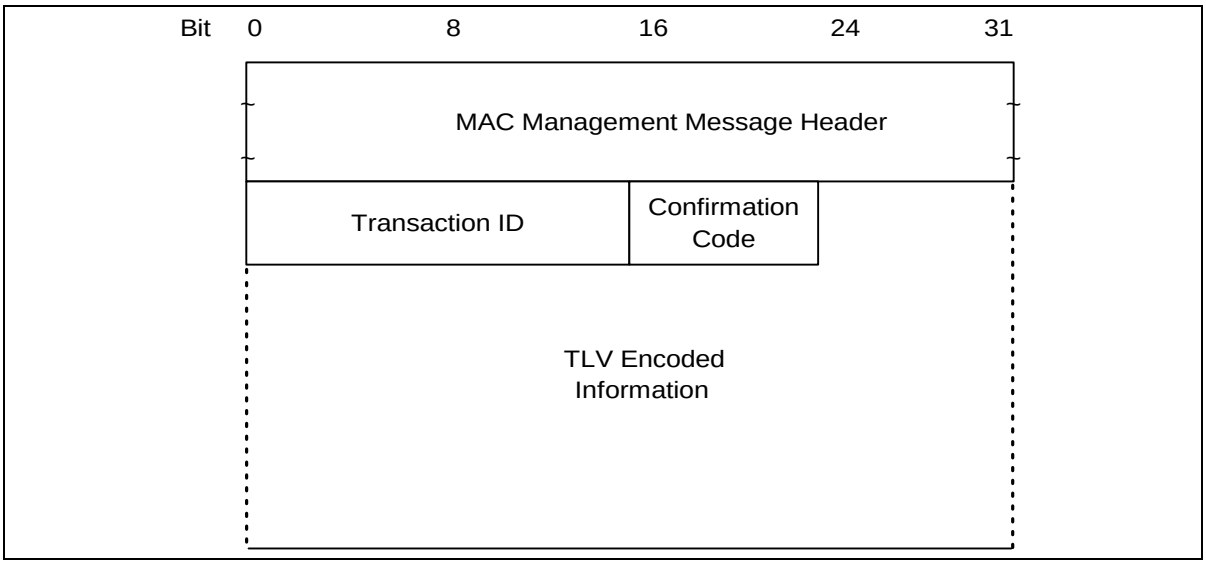


Figure 6-32. Dynamic Service Addition — Response

Parameters MUST be as follows:

Transaction ID Transaction ID from corresponding DSA-REQ.

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the entire corresponding DSA-Request.

All other parameters are coded as TLV tuples as defined in Appendix C.

If the transaction is successful, the DSA-RSP MAY contain one or more of the following:

Classifier Parameters The CMTS MUST include the complete specification of the Classifier in the DSA-RSP, including a newly assigned Classifier Identifier. The CM MUST NOT include the specification of the Classifier in the DSA-RSP.

Service Flow Parameters The CMTS MUST include the complete specification of the Service Flow in the DSA-RSP, including a newly assigned Service Flow Identifier and an expanded service class name if applicable. The CM MUST NOT include the specification of the Service Flow in the DSA-RSP.

Payload Header Suppression Parameters The CMTS MUST include the complete specification of the PHS Parameters in the DSA-RSP, including a newly assigned PHS Index, a Classifier Identifier and a Service Flow Identifier. The CM MUST NOT include the specification of the PHS Parameters.

If the transaction is unsuccessful due to Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, and the Confirmation Code is not one of the major error codes in Appendix C.4.2, the DSA-RSP MUST contain at least one of the following:

Service Flow Error Set A Service Flow Error Set and identifying Service Flow Reference/Identifier MUST be included for at least one failed Service Flow in the corresponding DSA-REQ. Every Service Flow Error Set MUST include at least one specific failed QoS Parameter of the corresponding Service Flow. This parameter MUST be omitted if the entire DSA-REQ is successful.

Classifier Error Set A Classifier Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for at least one failed Classifier in the corresponding DSA-REQ. Every Classifier Error Set MUST include at least one specific failed Classifier Parameter of the corresponding Classifier. This parameter MUST be omitted if the entire DSA-REQ is successful.

Payload Header Suppression Error Set A PHS Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for at least one failed PHS Rule in the corresponding DSA-REQ. Every PHS Error Set MUST include at least one specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter MUST be omitted if the entire DSA-REQ is successful.

If Privacy is enabled, the DSA-RSP message MUST contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.13.1 CM-Initiated Dynamic Service Addition

The CMTS's DSA-Response for Service Flows that are successfully added MUST contain a Service Flow ID. The DSA-Response for successfully Admitted or Active upstream QoS Parameter Sets MUST also contain a Service ID.

If the corresponding DSA-Request uses the Service Class Name (refer to Appendix C.2.2.3.4) to request service addition, a DSA-Response MUST contain the QoS Parameter Set associated with the named Service Class. If the Service Class Name is used in conjunction with other QoS Parameters in the DSA-Request, the CMTS MUST accept or reject the DSA-Request using the explicit QoS Parameters in the DSA-Request. If these Service Flow

Encodings conflict with the Service Class attributes, the CMTS MUST use the DSA-Request values as overrides for those of the Service Class.

If the transaction is successful, the CMTS MUST assign a Classifier Identifier to each requested Classifier and a PHS Index to each requested PHS Rule. The CMTS MUST use the original Classifier Reference(s) and Service Flow Reference(s) to link the successful parameters in the DSA-RSP.

If the transaction is unsuccessful, the CMTS MUST use the original Classifier Reference(s) and Service Flow Reference(s) to identify the failed parameters in the DSA-RSP.

6.3.13.2 CMTS-Initiated Dynamic Service Addition

If the transaction is unsuccessful, the CM MUST use the Classifier Identifier(s) and Service Flow Identifier(s) to identify the failed parameters in the DSA-RSP.

6.3.14 Dynamic Service Addition — Acknowledge (DSA-ACK)

A Dynamic Service Addition Acknowledge MUST be generated in response to a received DSA-RSP. The format of a DSA-ACK MUST be as shown in Figure 6-33.

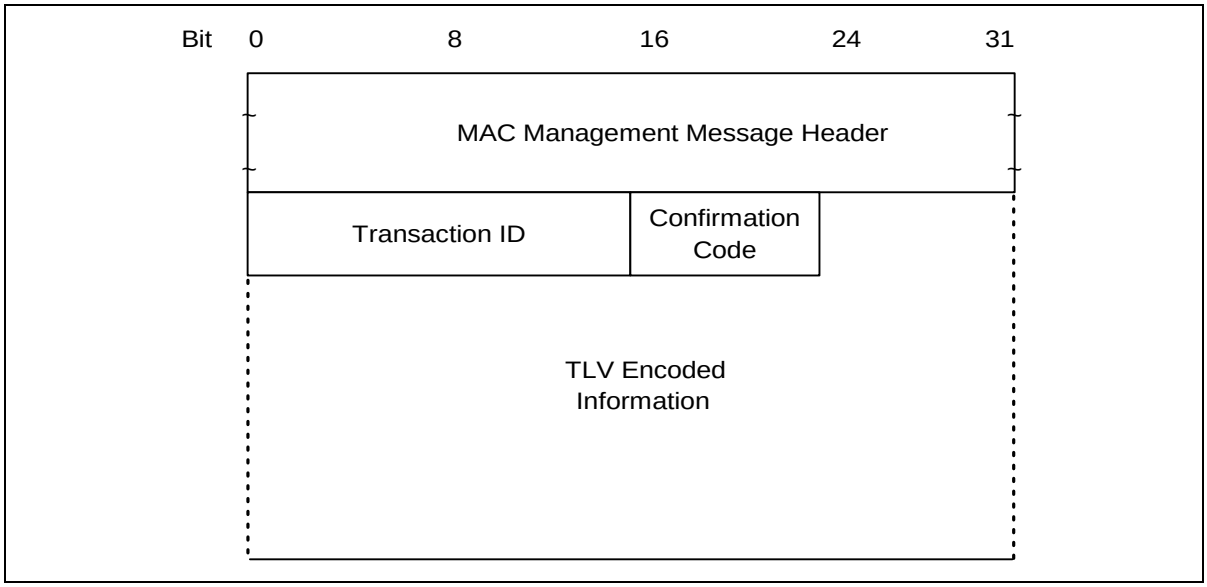


Figure 6-33. Dynamic Service Addition — Acknowledge

Parameters MUST be as follows:

Transaction ID Transaction ID from corresponding DSA-Response.

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the entire corresponding DSA-Response.[9]

9. The confirmation code is necessary particularly when a Service Class Name (refer to Section 8.1.3) is used in the DSA-Request. In this case, the DSA-Response could contain Service Flow parameters that the CM is unable to support (either temporarily or as configured).

All other parameters are coded TLV tuples.

Service Flow Error Set The Service Flow Error Set of the DSA-ACK message encodes specifics of failed Service Flows in the DSA-RSP message. A Service Flow Error Set and identifying Service Flow Reference/Identifier MUST be included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding DSA-REQ. This parameter MUST be omitted if the entire DSA-REQ is successful.

If Privacy is enabled, the DSA-RSP message MUST contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.15 Dynamic Service Change — Request (DSC-REQ)

A Dynamic Service Change Request MAY be sent by a CM or CMTS to dynamically change the parameters of an existing Service Flow. DSCs changing classifiers MUST carry the entire classifier TLV set for that new classifier.

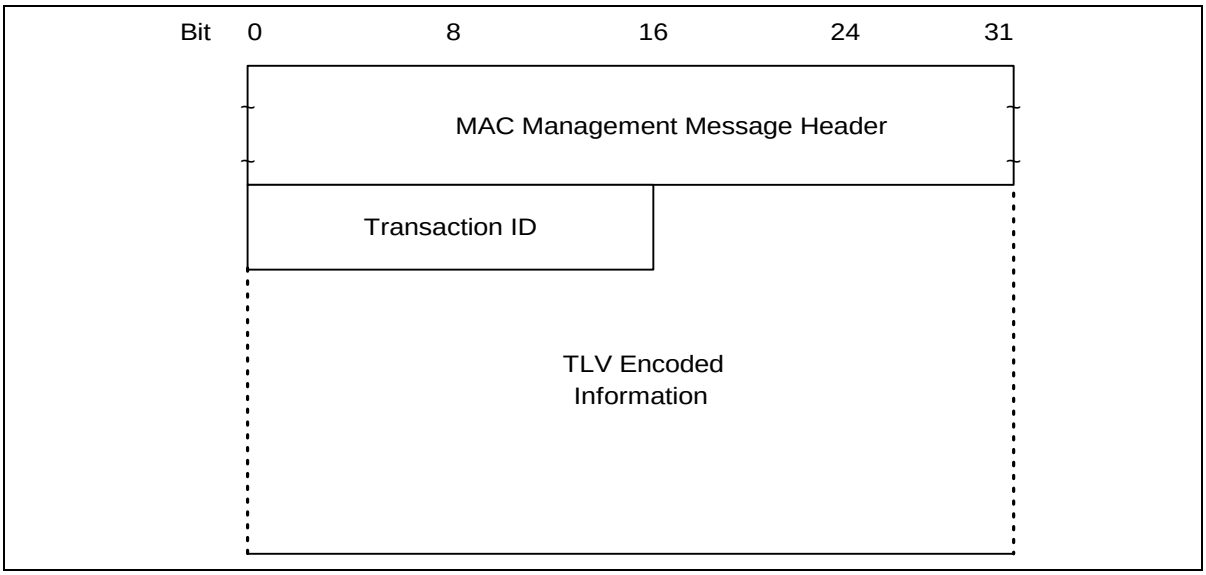


Figure 6-34. Dynamic Service Change — Request

A CM or CMTS MUST generate DSC-REQ messages in the form shown in Figure 6-34 including the following parameters:

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C. A DSC-REQ message MUST NOT carry parameters for more than one Service Flow in each direction, *i.e.*, a DSC-REQ message MUST contain

parameters for either a single upstream Service Flow, or for a single downstream Service Flow, or for one upstream and one downstream Service Flow. A DSC-REQ MUST contain at least one of the following:

Classifier Parameters Specification of the rules to be used to classify packets into a specific service flow - this includes the Dynamic Service Change Action TLV which indicates whether this Classifier should be added, replaced or deleted from the Service Flow (refer to Appendix C.2.1.3.7). If included, the Classifier Parameters MUST contain the Dynamic Change Action TLV, a Classifier Reference/ Identifier and a Service Flow Identifier.

Service Flow Parameters Specification of the Service Flow's new traffic characteristics and scheduling requirements. The Admitted and Active Quality of Service Parameter Sets in this message replace the Admitted and Active Quality of Service Parameter Sets currently in use by the Service Flow. If the DSC message is successful and it contains Service Flow parameters, but does not contain replacement sets for both Admitted and Active Quality of Service Parameter Sets, the omitted set(s) MUST be set to null. If included, the Service Flow Parameters MUST contain a Service Flow Identifier.

Payload Header Suppression Parameters Specification of the rules to be used for Payload Header Suppression to suppress payload headers related to a specific Classifier—this includes the Dynamic Service Change Action TLV which indicates whether this PHS Rule should be added, set or deleted from the Service Flow or whether all the PHS Rules for the Service Flow specified should be deleted (refer to Appendix C.2.2.8.5). If included, the PHS parameters MUST contain the Dynamic Service Change Action TLV, a Classifier Reference/Identifier, and a Service Flow Identifier, unless the Dynamic Service Change Action is "Delete all PHS Rules".[2] If the Dynamic Service Change Action is "Delete all PHS Rules", the PHS Parameters MUST contain a Service Flow Identifier along with the Dynamic Service Change Action, and no other PHS parameters need be present in this case. However, if other PHS parameters are present, in particular Payload Header Suppression Index, they MUST be ignored by the receiver of the DSC-REQ message.

If Privacy is enabled, a DSC-REQ MUST also contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.16 Dynamic Service Change — Response (DSC-RSP)

A Dynamic Service Change Response MUST be generated in response to a received DSC-REQ. The format of a DSC-RSP MUST be as shown in Figure 6-35

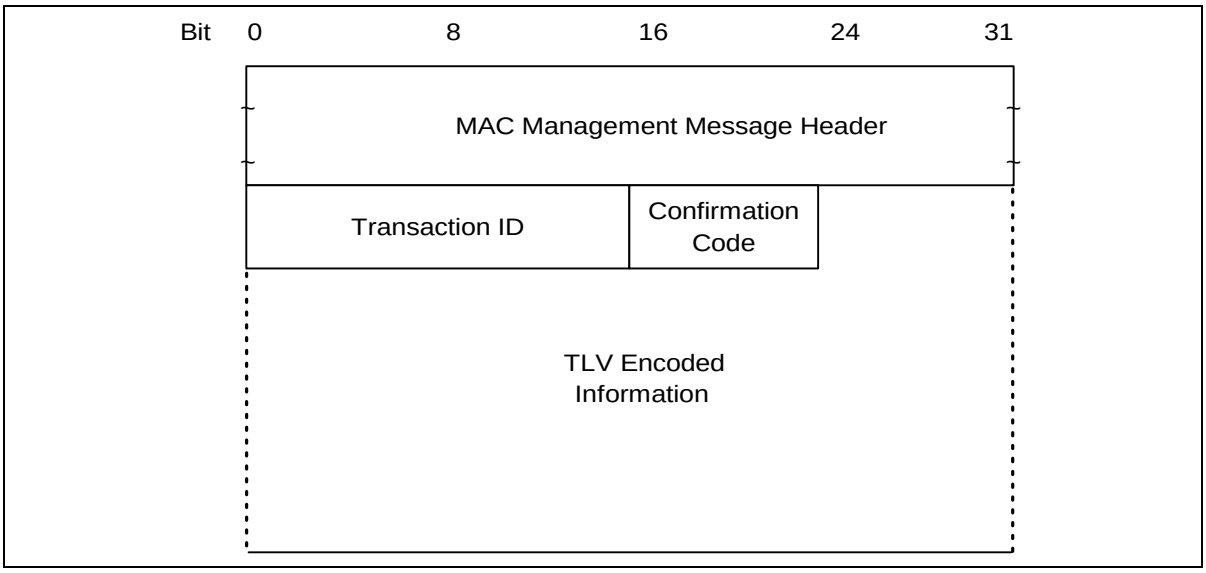


Figure 6-35. Dynamic Service Change — Response

Parameters MUST be as follows:

Transaction ID Transaction ID from corresponding DSC-REQ

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the corresponding DSC-Request.

All other parameters are coded as TLV tuples as defined in Appendix C.

If the transaction is successful, the DSC-RSP MAY contain one or more of the following:

Classifier Parameters The CMTS MUST include the complete specification of the Classifier in the DSC-RSP, including a newly assigned Classifier Identifier for new Classifiers. The CM MUST NOT include the specification of the Classifier in the DSC-RSP.

Service Flow Parameters The CMTS MUST include the complete specification of the Service Flow in the DSC-RSP, including an expanded service class name if applicable. The CMTS MUST include a SID in the DSC-RSP if a Service Flow Parameter Set contained an upstream Admitted QoS Parameter Set and this Service Flow does not have an associated SID. If a Service Flow Parameter set contained a Service Class Name and an Admitted QoS Parameter Set, the CMTS MUST include the QoS Parameter Set corresponding to the named Service Class in the DSC-RSP. If specific QoS Parameters were also included in the classed Service Flow request, the CMTS MUST include these QoS Parameters in the DSC-RSP instead of any QoS Parameters of the same type of the named Service Class. The CM MUST NOT include the specification of the Service Flow in the DSC-RSP.

Payload Header Suppression Parameters — The CMTS MUST include the complete specification of the PHS Parameters in the DSC-RSP, including a newly assigned PHS Index for new PHS rules, a Classifier Identifier and a Service Flow Identifier. The CM MUST NOT include the specification of the PHS Parameters.

If the transaction is unsuccessful due to Service Flow Parameters, Classifier Parameters, or Payload Header Suppression Parameters, and the Confirmation Code is not one of the major error codes in Appendix C.4.2, the DSC-RSP MUST contain at least one of the following:

Classifier Error Set — A Classifier Error Set and identifying Classifier Reference/Identifier and Service Flow Reference/Identifier pair MUST be included for at least one failed Classifier in the corresponding DSC-REQ. Every Classifier Error Set MUST include at least one specific failed Classifier Parameter of the corresponding Classifier. This parameter MUST be omitted if the entire DSC-REQ is successful.

Service Flow Error Set — A Service Flow Error Set and identifying Service Flow ID MUST be included for at least one failed Service Flow in the corresponding DSC-REQ. Every Service Flow Error Set MUST include at least one specific failed QoS Parameter of the corresponding Service Flow. This parameter MUST be omitted if the entire DSC-REQ is successful.

Payload Header Suppression Error Set — A PHS Error Set and identifying Service Flow Reference/Identifier and Classifier Reference/Identifier pair MUST be included for at least one failed PHS Rule in the corresponding DSC-REQ, unless the Dynamic Service Change Action is “Delete all PHS Rules.” If the Dynamic Service Change Action is “Delete all PHS Rules” the PHS Error Set(s) MUST include an identifying Service Flow ID. Every PHS Error Set MUST include at least one specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter MUST be omitted if the entire DSC-REQ is successful.”

Regardless of success or failure, if Privacy is enabled for the CM the DSC-RSP MUST contain:

Key Sequence Number — The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest — The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message’s Attribute list. (Refer to Appendix C.1.4.1)

6.3.17 Dynamic Service Change — Acknowledge (DSC-ACK)

A Dynamic Service Change Acknowledge MUST be generated in response to a received DSC-RSP. The format of a DSC-ACK MUST be as shown in Figure 6-36

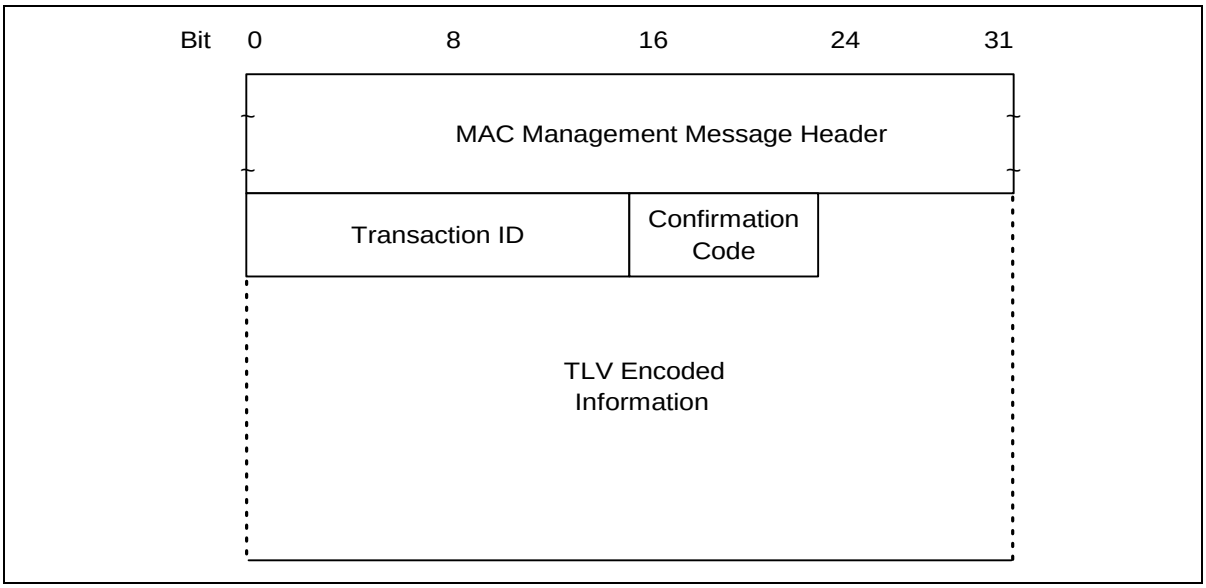


Figure 6-36. Dynamic Service Change — Acknowledge

Parameters MUST be as follows:

Transaction ID Transaction ID from the corresponding DSC-REQ

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the entire corresponding DSC-Response.[10]

All other parameters are coded TLV tuples.

Service Flow Error Set The Service Flow Error Set of the DSC-ACK message encodes specifics of failed Service Flows in the DSC-RSP message. A Service Flow Error Set and identifying Service Flow Identifier MUST be included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding DSC-REQ. This parameter MUST be omitted if the entire DSC-REQ is successful.

If Privacy is enabled, the DSC-ACK message MUST contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

[10] The Confirmation Code and Service Flow Error Set are necessary particularly when a Service Class Name is (refer to Section 8.1.3) used in the DSC-Request. In this case, the DSC-Response could contain Service Flow parameters that the CM is unable to support (either temporarily or as configured).

6.3.18 Dynamic Service Deletion — Request (DSD-REQ)

A DSD-Request MAY be sent by a CM or CMTS to delete a single existing Upstream Service Flow and/or a single existing Downstream Service Flow. The format of a DSD-Request MUST be as shown in Figure 6-37.

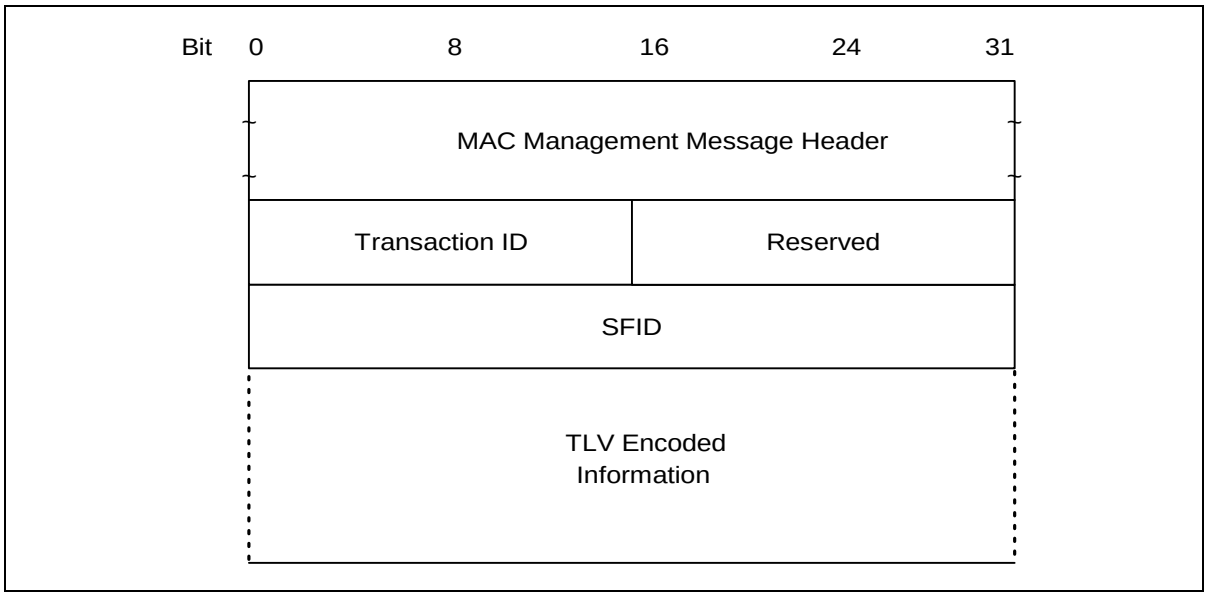


Figure 6-37. Dynamic Service Deletion — Request

Parameters MUST be as follows:

Service Flow Identifier If this value is non-zero it is the SFID of a single Upstream or single Downstream Service Flow to be deleted. If this value is zero the Service Flow(s) to be deleted will be identified by SFID(s) in the TLV's. If this value is non-zero then any SFIDs included in the TLV's MUST be ignored.

Transaction ID Unique identifier for this transaction assigned by the sender.

All other parameters are coded as TLV tuples as defined in Appendix C.

Service Flow Identifier The SFID(s) to be deleted, which MUST be encoded per Appendix C.2.2.3.2 The Service Flow Identifier TLV MUST be the only Service Flow Encoding sub-TLV used.

If Privacy is enabled, the DSD-REQ MUST include:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.19 Dynamic Service Deletion – Response (DSD-RSP)

A DSD-RSP MUST be generated in response to a received DSD-REQ. The format of a DSD-RSP MUST be as shown in Figure 6-38.

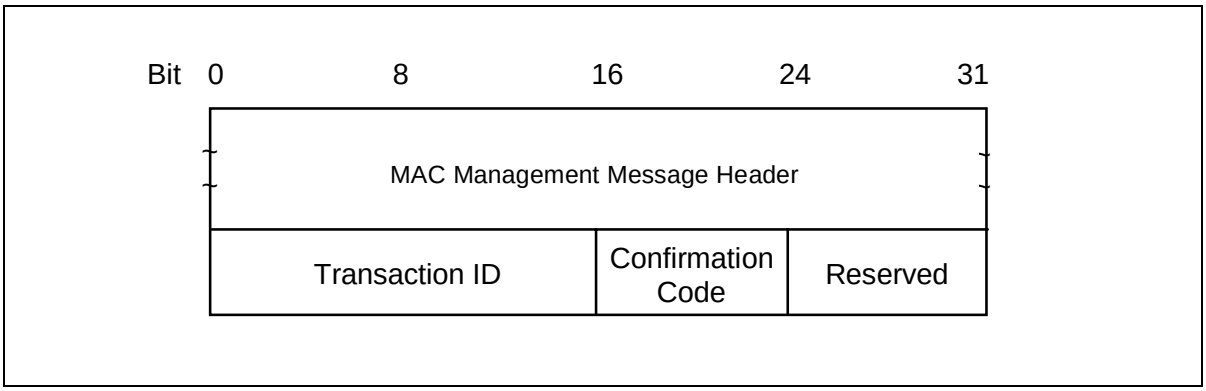


Figure 6-38. Dynamic Service Deletion — Response

Parameters MUST be as follows:

Transaction ID Transaction ID from corresponding DSD-REQ.

Confirmation Code The appropriate Confirmation Code (refer to Appendix C.4) for the corresponding DSD-Request.

6.3.20 Dynamic Channel Change – Request (DCC-REQ)

A Dynamic Channel Change Request MAY be transmitted by a CMTS to cause a CM to change the upstream channel on which it is transmitting, the downstream channel it is receiving, or both.

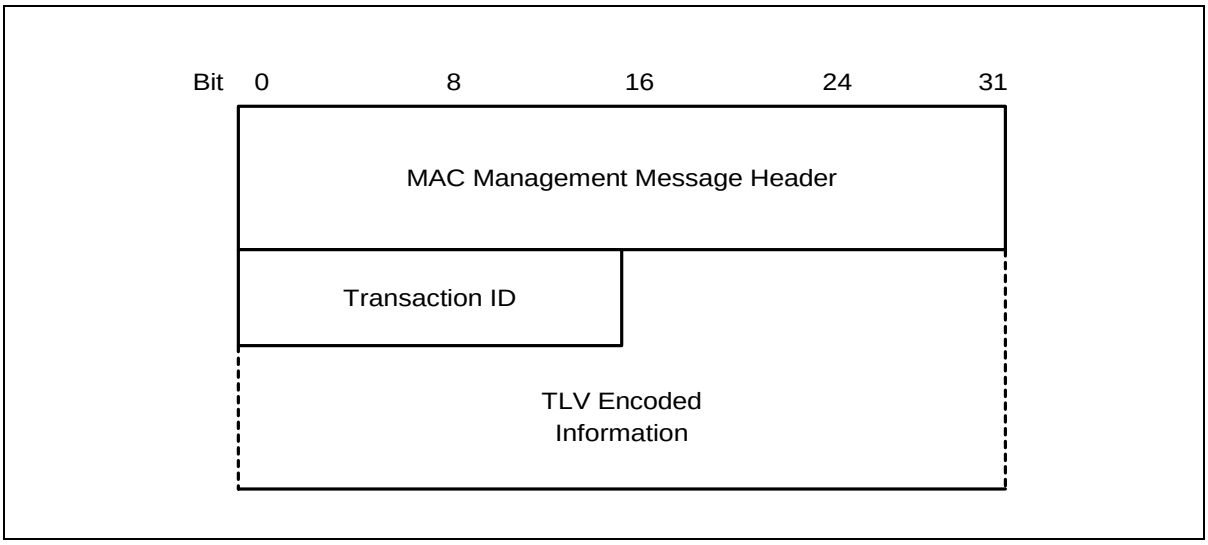


Figure 6-39. Dynamic Channel Change Request

A CMTS MUST generate DCC-REQ message in the form shown in Figure 6-39 including the following parameter:

Transaction ID A 16-bit unique identifier for this transaction assigned by the sender.

The following parameters are optional and are coded as TLV tuples:

Upstream Channel ID The identifier of the upstream channel to which the CM is to switch for upstream transmissions.

Downstream Parameters The frequency of the downstream channel to which the CM is to switch for downstream reception.

Initialization Technique Directions for the type of initialization, if any, that the CM should perform once synchronized to the new channel(s).

UCD Substitution Provides a copy of the UCD for the new channel.

SAID Substitution A pair of Security Association Identifiers (SAID) which contain the current SAID and the new SAID for the new channel. This TLV occurs once if the SAID requires substitution.

Service Flow Substitution A group of sub-TLVs which allows substitution in a Service Flow of the Service Flow Identifier, Service Identifier, Classifier Identifier, and the Payload Header Suppression Index. This TLV is repeated for every Service Flow which has parameters requiring substitution.

If Privacy is enabled, a DCC-REQ MUST also contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Channel Change message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.20.1 Encodings

The type values used MUST be those shown below. These are unique within the Dynamic Channel Change Request message, but not across the entire MAC message set.

If a CM performs a channel change without performing a re-initialization (as defined in Section 6.3.20.1.3), then all the configuration variables of the CM MUST remain constant, with the exception of the configuration variables which are explicitly changed below. The CM will not be aware of any configuration changes other than the ones that have been supplied in the DCC command, so consistency in provisioning between the old and new channels is important.

6.3.20.1.1 Upstream Channel ID

When present, this TLV specifies the new upstream channel ID that the CM MUST use when performing a Dynamic Channel Change. It is an override for the current upstream channel ID. The CMTS MUST ensure that the Upstream Channel ID for the new channel is different than the Upstream Channel ID for the old channel. This TLV MUST be included if the upstream channel is changed, even if the UCD substitution TLV is included.

Type	Length	Value
1	1	0-255: Upstream Channel ID

If this TLV is missing, the CM MUST NOT change its upstream channel ID. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.2 Downstream Parameters

When present, this TLV specifies the operating parameters of the new downstream channel. The value field of this TLV contains a series of sub-types.

Type	Length	Value
2	n	

The CMTS MAY include this TLV. If this TLV is missing, the CM MUST NOT change its downstream parameters.

6.3.20.1.2.1 Downstream Frequency

This TLV specifies the new receive frequency that the CM MUST use when performing a Dynamic Channel Change. It is an override for the current downstream channel frequency. This is the center frequency of the

downstream channel in Hz and is stored as a 32-bit binary number. The downstream frequency MUST be a multiple of 62,500 Hz.

Subtype	Length	Value
2.1	4	Rx Frequency

The CMTS MUST include this sub-TLV. The CM MUST observe this sub-TLV.

6.3.20.1.2.2 Downstream Modulation Type

This TLV specifies the modulation type that is used on the new downstream channel.

Subtype	Length	Value
2.2	1	0 = 64 QAM 1 = 256 QAM 2 - 255: reserved

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

6.3.20.1.2.3 Downstream Symbol Rate

This TLV specifies the symbol rate that is used on the new downstream channel.

Subtype	Length	Value
2.3	1	0 = 5.056941 Msym/sec 1 = 5.360537 Msym/sec 2 = 6.952 Msym/sec 3 - 255: reserved

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

6.3.20.1.2.4 Downstream Interleaver Depth

This TLV specifies the parameters “I” and J of the downstream interleaver.

Subtype	Length	Value
2.4	2	I: 0-255 J: 0-255

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

6.3.20.1.2.5 Downstream Channel Identifier

This TLV specifies the 8 bit downstream channel identifier of the new downstream channel.

Subtype	Length	Value
2.5	1	0-255: Downstream Channel ID.

The CMTS SHOULD include this sub-TLV. The CM SHOULD observe this sub-TLV.

6.3.20.1.2.6 SYNC Substitution

When present, this TLV allows the CMTS to inform the CM to wait or not wait for a SYNC message before proceeding. The CMTS MUST have synchronized timestamps between the old and new channel(s) if it instructs the CM to not wait for a SYNC message before transmitting on the new channel. Synchronized timestamps implies that the timestamps are derived from the same clock and contain the same value.

Type	Length	Value
2.6	1	0 = acquire SYNC message on the new downstream channel before proceeding 1 = proceed without first obtaining the SYNC message 2 - 255: reserved

If this TLV is absent, the CM MUST wait for a SYNC message on the new channel before proceeding. If the CM has to wait for a new SYNC message when changing channels, then operation may be suspended for a time up to the “SYNC Interval” (see Appendix B) or longer, if the SYNC message is lost or is not synchronized with the old channel(s).

An alternative approach is to send SYNC messages more frequently (every 10 ms for example), and continue to require the CM to wait for a SYNC message before proceeding. This approach has slightly more latency, but provides an additional check to prevent the CM from transmitting at an incorrect time interval.

The CMTS SHOULD include this TLV. The CM SHOULD observe this TLV.

6.3.20.1.3 Initialization Technique

When present, this TLV allows the CMTS to direct the CM as to what level of re-initialization, if any, it MUST perform before it can commence communications on the new channel(s). The CMTS can make this decision based upon its knowledge of the differences between the old and new MAC domains and the PHY characteristics of their upstream and downstream channels.

Typically, if the move is between upstream and/or downstream channels within the same MAC domain, then the connection profile values may be left intact. If the move is between different MAC domains, then a complete initialization may be performed.

If a complete re-initialization is not required, some re-ranging may still be required. For example, areas of upstream spectrum are often configured in groups. A DCC-REQ to an adjacent upstream channel within a group may not warrant re-ranging. Alternatively, a DCC-REQ to a non-adjacent upstream channel might require station maintenance whereas a DCC-REQ from one upstream channel group to another might require initial

maintenance. Re-ranging may also be required if there is any difference in the PHY parameters between the old and new channels.

Type	Length	Value
3	1	0 = Reinitialize the MAC 1 = Perform initial maintenance on new channel before normal operation 2 = Perform station maintenance on new channel before normal operation 3 = Perform either initial maintenance or station maintenance on new channel before normal operation 4 = Use the new channel(s) directly without re-initializing or performing initial or station maintenance 5 - 255: reserved

The CM MUST first select the new upstream and downstream channels based upon the Upstream Channel ID TLV (refer to Section 6.3.20.1.1) and the Downstream Frequency TLV (refer to Section 6.3.20.1.2.1). Then the CM MUST follow the directives of this TLV. For option 0, the CM MUST begin with the Initialization SID. For options 1 to 4 the CM MUST continue to use the primary SID for ranging. A SID Substitution TLV (see Section 6.3.21.1.7.2) may specify a new primary SID for use on the new channel (refer to Section 6.3.20.1.6.2).

Option 0: This option directs the CM to perform all the operations associated with initializing the CM (refer to Section 9.2). This includes all the events after acquiring downstream QAM, FEC, and MPEG lock and before Standard Operation (refer to Section 9.3), including obtaining a UCD, ranging, establishing IP connectivity, establishing time of day, transfer of operational parameters, registration, and baseline privacy initialization. When this option is used, the only other TLVs in DCC-REQ that are relevant are the Upstream Channel ID TLV and the Downstream Parameters TLV. All other DCC-REQ TLVs are irrelevant.

Option 1: If initial maintenance is specified, operation on the new channel could be delayed by several Ranging Intervals (see Appendix B).

Option 2: If station maintenance is specified, operation on the new channel could be delayed by the value of T4 (see Appendix B).

Option 3: This value authorizes a CM to use an initial maintenance or station maintenance region, which ever the CM selects. This value might be used when there is uncertainty when the CM may execute the DCC command and thus a chance that it might miss station maintenance slots.

Option 4: This option provides for the least interruption of service, and the CM may continue its normal operation as soon as it has achieved synchronization on the new channel. This option is intended for use with a near-seamless channel change (refer to Section 9.4.5.4).

Note: This option should not be used in physical plants where upstream transmission characteristics are not consistent.

If this TLV is absent, the CM MUST re-initialize the MAC. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.4 UCD Substitution

When present, this TLV allows the CMTS to send an Upstream Channel Descriptor message to the CM. This UCD message is intended to be associated with the new upstream and/or downstream channel(s). The CM stores this UCD message in its cache, and uses it after synchronizing to the new channel(s).

Type	Length	Value
4	n	UCD for the new upstream channel

This TLV includes all parameters for the UCD message as described in section Section 6.3.3 except for the MAC Management Message Header. The CMTS MUST ensure that the change count in the UCD matches the change count in the UCDs of the new channel(s). The CMTS MUST ensure that the Upstream Channel ID for the new channel is different than the Upstream Channel ID for the old channel.

If the length of the UCD exceeds 254 bytes, the UCD MUST be fragmented into two or more successive Type 4 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the UCD Substitution by concatenating the contents (Value of the TLV) of successive Type 4 elements in the order in which they appear in the DCC-REQ message. For example, the first byte following the length field of the second Type 4 element is treated as if it immediately follows the last byte of the first Type 4 element.

If the CM has to wait for a new UCD message when changing channels, then operation may be suspended for a time up to the “UCD Interval” (see Appendix B) or longer, if the UCD message is lost.

The CMTS SHOULD include this TLV. The CM SHOULD observe this TLV.

6.3.20.1.5 Security Association Identifier (SAID) Substitution

When present, this TLV allows the CMTS to replace the Security Association Identifier (SAID) in the current Service Flow with a new Security Association Identifier. The baseline privacy keys associated with the SAID MUST remain the same. The CM does not have to simultaneously respond to the old and new SAID.

Type	Length	Value
6	4	current SAID (lower order 14 bits of a 16 bits field), new SAID (lower order 14 bits of a 16 bit field)

If this TLV is absent, the current Security Association Identifier assignment is retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.6 Service Flow Substitutions

When present, this TLV allows the CMTS to replace specific parameters within the current Service Flows on the current channel assignment with new parameters for the new channel assignment. One TLV is used for each Service Flow that requires changes in parameters. The CMTS may choose to do this to help facilitate setting up new QoS reservations on the new channel before deleting QoS reservations on the old channel. The CM does not have to simultaneously respond to the old and new Service Flows.

This TLV allows resource assignments and services to be moved between two independent ID value spaces and scheduling entities by changing the associated IDs and indexes. ID value spaces that may differ between the two channels include the Service Flow Identifier and the Service ID. This TLV does not allow changes to Service Flow QoS parameters.

The Service Class Names used within the Service Flow ID should remain identical between the old and new channels.

Type	Length	Value
7	n	list of subtypes

If this TLV is absent for a particular Service Flow, then current Service Flow and its attributes are retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.6.1 Service Flow Identifier Substitution

This TLV allows the CMTS to replace the current Service Flow Identifier (SFID) with a new Service Flow Identifier. Refer to Appendix C.2.2.3.2 for details on the usage of this parameter.

This TLV MUST be present if any other Service Flow subtype substitutions are made. If this TLV is included and the Service Flow ID is not changing, then the current and new Service Flow ID will be set to the same value.

Subtype	Length	Value
7.1	8	current Service Flow ID, new Service Flow ID

The CMTS MUST include this Sub-TLV. The CM MUST observe this Sub-TLV.

6.3.20.1.6.2 Service Identifier Substitution

When present, this TLV allows the CMTS to replace the Service Identifier (SID) in the current upstream Service Flow with a new Service Identifier. Refer to Appendix C.2.2.3.3 for details on the usage of this parameter.

Subtype	Length	Value
7.2	4	current SID (lower order 14 bits of a 16 bits field), new SID (lower order 14 bits of a 16 bits field)

If this TLV is absent, the current Service Identifier assignments are retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.6.3 Unsolicited Grant Time Reference Substitution

When present, this TLV allows the CMTS to replace the current Unsolicited Grant Time Reference with a new Unsolicited Grant Time Reference. Refer to Appendix C.2.2.6.11 for details on the usage of this parameter.

This TLV is useful if the old and new upstream use different time bases for their time stamps. This TLV is also useful if the Unsolicited Grant transmission window is moved to a different point in time. Changing this value may cause operation to temporarily exceed the jitter window specified by Appendix C.2.2.6.8.

Subtype	Length	Value
7.5	4	new reference

If this TLV is absent, the current Unsolicited Grant Time Reference is retained. The CMTS MAY include this TLV. The CM MUST observe this TLV.

6.3.20.1.7 CMTS MAC Address

When present, this TLV allows the current CMTS to send the MAC address of the destination CMTS corresponding to the target downstream frequency. This TLV MUST be specified if the CM is changing downstream channels and UCD substitution is specified or if the CM is changing downstream channels and using initialization technique 4, use the new channel(s) directly.

Type	Length	Value
8	6	MAC Address of Destination CMTS

The CMTS SHOULD include the CMTS MAC address TLV. The CM SHOULD observe the CMTS MAC address TLV.

6.3.21 Dynamic Channel Change – Response (DCC-RSP)

A Dynamic Channel Change Response MUST be transmitted by a CM in response to a received Dynamic Channel Change Request message to indicate that it has received and is complying with the DCC-REQ. The format of a DCC-RSP message MUST be as shown in Figure 6-40.

Before it begins to switch to a new upstream or downstream channel, a CM MUST transmit a DCC-RSP on its existing upstream channel. When a CM receives a DCC-REQ message requesting that it switch to an upstream and downstream channel that it is already using or requesting that it switch to only an upstream or downstream channel that it is already using, the CM MUST respond with a DCC-RSP message on that channel indicating that it is already using the correct channel.

A CM MAY ignore a DCC-REQ message while it is in the process of performing a channel change.

After switching to a new channel, if the MAC was not re-initialized per DCC-REQ Initialization TLV, option 0, the CM MUST send a DCC-RSP message to the CMTS. A DCC-RSP MUST NOT be sent if the CM reinitializes its MAC.

The full procedure for changing channels is described in Section 9.4.5.

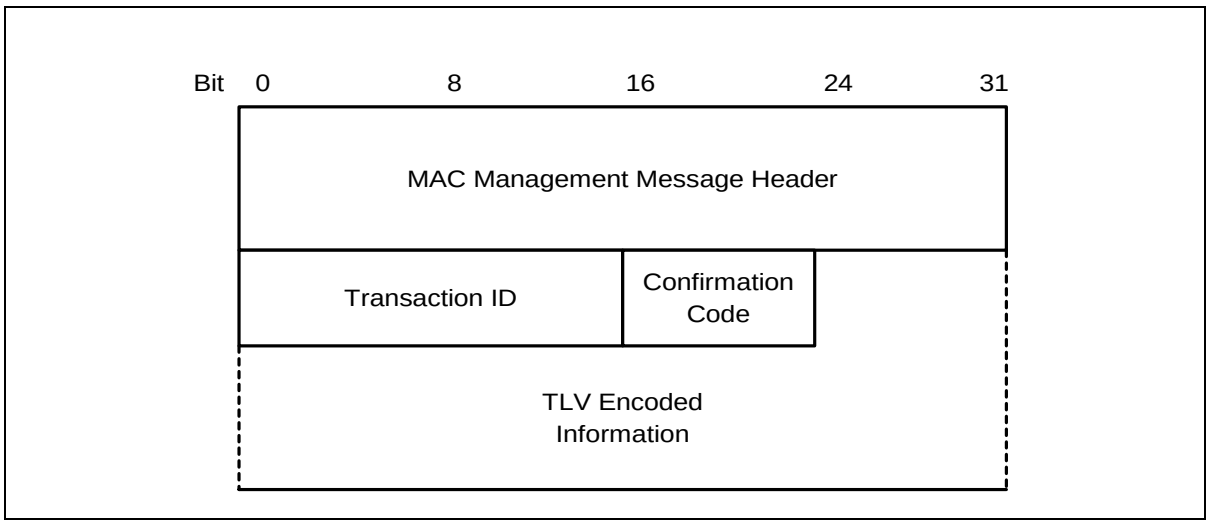


Figure 6-40. Dynamic Channel Change Response

Parameters MUST be as follows:

Transaction ID A 16 bit Transaction ID from corresponding DCC-REQ

Confirmation Code An 8 bit Confirmation Code as described in Appendix C.4.1

The following parameters are optional and are coded as TLV tuples.

CM Jump Time Timing parameters describing when the CM will make the jump

Regardless of success or failure, if Privacy is enabled for the CM the DCC-RSP MUST contain:

Key Sequence Number	The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest (refer to Appendix C.1.4.3)
HMAC-Digest	The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Channel Change message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.21.1 Encodings

The type values used MUST be those shown below. These are unique within the Dynamic Channel Change Response message, but not across the entire MAC message set.

6.3.21.1.1 CM Jump Time

When present, this TLV allows the CM to indicate to the CMTS when the CM plans to perform its jump and be disconnected from the network. With this information, the CMTS MAY take preventative measures to minimize or to eliminate packet drops in the downstream due to the channel change.

Type	Length	Value
1	n	

The time reference and units of time for these sub-TLVs is based upon the same 32-bit time base used in the SYNC message on the current downstream channel. This timestamp is incremented by a 10.24 MHz clock

The CM SHOULD include this TLV. The CMTS SHOULD observe this TLV.

6.3.21.1.1.1 Length of Jump

This TLV indicates to the CMTS the length of the jump from the previous channel to the new channel. Specifically, it represents the length of time that the CM will not be able to receive data in the downstream.

Subtype	Length	Value
1	4	length (based upon timestamp)

The CM MUST include this sub-TLV.

6.3.21.1.1.2 Start Time of Jump

When present, this TLV indicates to the CMTS the time in the future that the CM is planning on making the jump.

Subtype	Length	Value
2	8	start time (based upon timestamp), accuracy of start time (based upon timestamp)

The 32 bit, 10.24 MHz time base rolls over approximately every 7 minutes. If the value of the start time is less than the current timestamp, the CMTS will assume one roll-over of the timestamp counter has elapsed. The accuracy of the start time is an absolute amount of time before and after the start time.

The potential jump window is from (start time - accuracy) to (start time + accuracy + length).

The CM SHOULD include this TLV.

6.3.22 Dynamic Channel Change – Acknowledge (DCC-ACK)

A Dynamic Channel Change Acknowledge MUST be transmitted by a CMTS in response to a received Dynamic Channel Change Response message on the new channel with its Confirmation Code set to arrive(1). The format of a DCC-ACK message MUST be as shown in Figure 6-41.

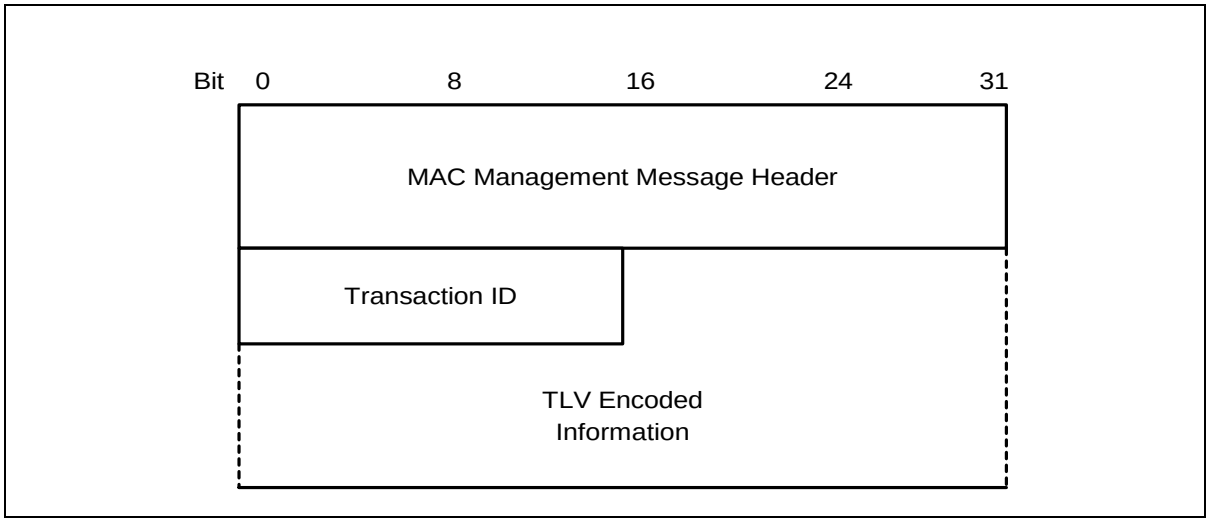


Figure 6-41. Dynamic Channel Change Acknowledge

Parameters MUST be as follows:

Transaction ID A 16 bit Transaction ID from corresponding DCC-RSP.

If Privacy is enabled, the DCC-ACK message MUST contain:

Key Sequence Number The key sequence number of the Auth Key, which is used to calculate the HMAC-Digest. (Refer to Appendix C.1.4.3)

HMAC-Digest The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Channel Change message's Attribute list. (Refer to Appendix C.1.4.1)

6.3.23 Device Class Identification Request (DCI-REQ)

A CM MAY implement the DCI-REQ message.

When implemented, a CM MUST transmit a DCI-REQ immediately following receipt of a ranging complete indication from the CMTS. A CM MUST NOT continue with initialization until a DCI-RSP message is received from the CMTS. Timeout and retry information is provided in Appendix B.

The DCI-REQ MUST be formatted as shown in Figure 6-42.

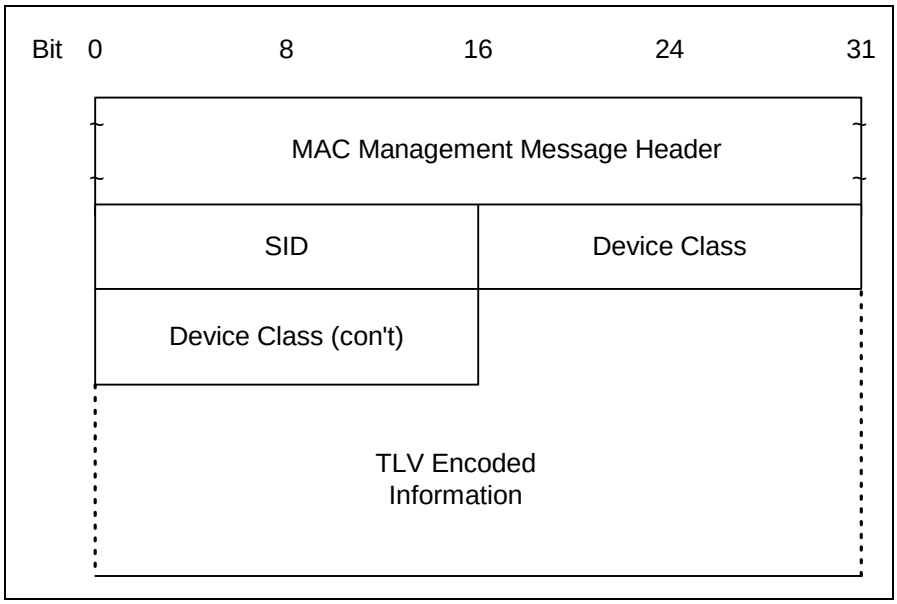


Figure 6-42. Device Class Identification Request

Parameters MUST be as follows:

SID: The temporary SID assigned during Ranging

Device Class: This is a 32-bit field where individual bits represent individual attributes of the CM. Bit #0 is the LSB of the field. Bits are set to 1 to select the attributes defined below.

bit #0 CPE Controlled Cable Modem (CCCM)

bits #1-31 reserved and must be set to zero

6.3.24 Device Class Identification Response (DCI-RSP)

A DCI-RSP MUST be transmitted by a CMTS in response to a received DCI-REQ.

The DCI-REQ MUST be formatted as shown in Figure 6-43.

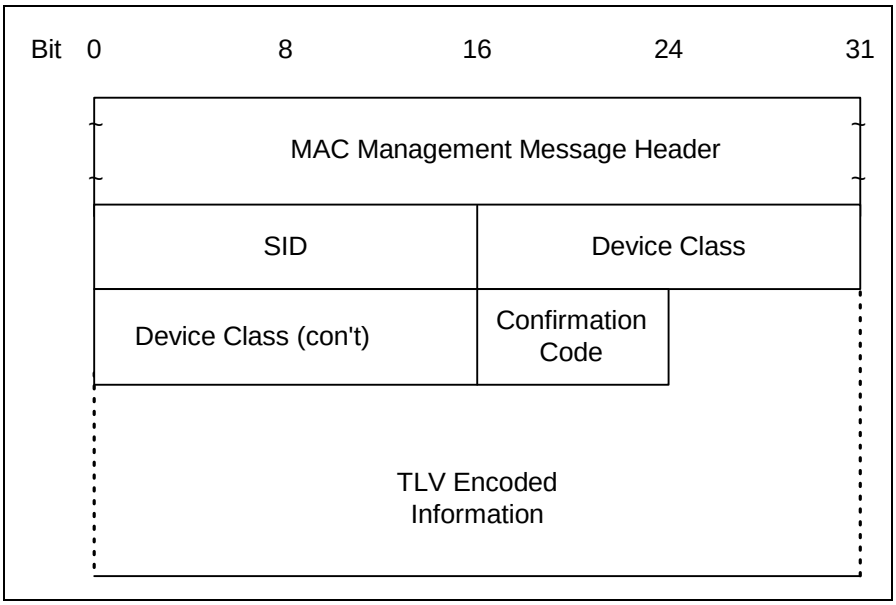


Figure 6-43. Device Class Identification Response

Parameters MUST be as follows:

SID: The SID received in the associated DCI-REQ

Device Class: The device class field as received in the associated DCI-REQ

Confirmation Code (refer to Appendix C.4):

The CMTS MUST use only one of 3 confirmation codes in the DCI-RSP.

- If the response is reject-temporary(3), the CM MUST reset its DCI-REQ retry counter to zero and MUST resend the DCI-REQ and wait for the DCI-RSP before proceeding.
- If the response is reject-permanent(4), the CM MUST abort this registration attempt and MUST begin rescanning for a different downstream channel. The CM MUST NOT retry this channel until it has tried all other DOCSIS downstream channels on the network.
- If the response is success(0), the CM MUST continue with registration.

The CMTS MUST retain the device class information for use in the DHCP Process. The CMTS MUST create a DHCP Agent Option 82 tuple with the device class information as specified in [RFC-3256] and MUST insert this tuple in the DHCPDISCOVER from the corresponding CM before forwarding that DHCPDISCOVER to the DHCP server.

6.3.25 Upstream Transmitter Disable (UP-DIS) MAC Management Message

The UP-DIS message provides additional functionality to permanently or temporarily disable the modem, as well as to disable the modem for a specified period of time. It is used to control the admission of certain modem types and groups to the network as early as immediately before registration. It can also be used for network

troubleshooting, disabling the modems that violate network policies, or for avoiding request floods in a large network, when the CMTS goes on-line.

This message is stateless and can be issued by the CMTS at any time. The UP-DIS message is sent from a CMTS to a CM; there is no response from the CM transmitted back to the CMTS. UP-DIS messages may be unicast, in which case the destination address in the MAC header is the address of the selected CM, or multicast, in which case the destination address is a well-known MAC multicast address (see Appendix A for details on well-known addresses).

The CMTS MUST be capable of transmitting the UP-DIS message. The CMTS can transmit UP-DIS messages either as a result of a triggering event detected by CMTS internally, or in response to a remote management command. Mechanisms for setting up, detecting, and reporting situations where the transmission of an UP-DIS message might be appropriate, are implementation dependent. Similarly, signaling, which remotely instructs CMTS to transmit of the UP-DIS message, is outside the scope of this specification. One of the possible implementations may be SNMP command sent to CMTS over network.

CMs SHOULD support the UP-DIS message for easier network management.

Since the UP-DIS mechanism at the CM is stateless and the CMs do not retain disabled status after power cycle, the CMTS MAY incorporate mechanisms to track disabled CMs by their MAC addresses. CMTS would resend an UP-DIS message as appropriate to the modems that were permanently disabled by the network operator, and then power cycled by the user to attempt to re-register. However, the same function may also be implemented by provisioning infrastructure on modem registration, and therefore if CMTS is unable to track disabled modems autonomously, it SHOULD be able to send UP-DIS in response to external command.

The UP-DIS message MUST be formatted as shown in Figure 6-44.

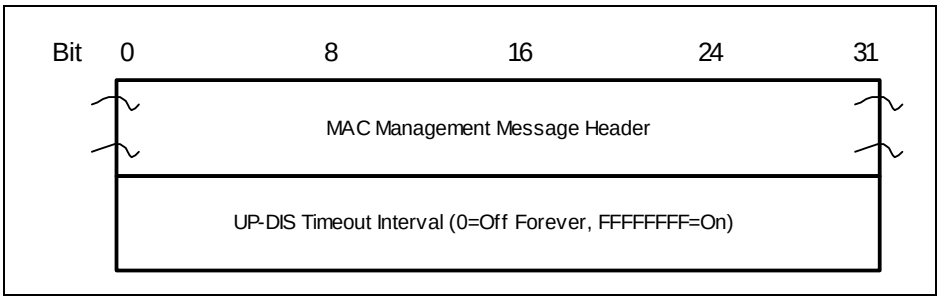


Figure 6-44. UP-DIS message format

The only parameter is UP-DIS Timeout Interval, which MUST be encoded as follows.

UP-DIS Timeout Interval: A 32-bit, unsigned integer representing the disable timeout interval in milliseconds. There are two special values defined:

00000000 permanently disables the upstream of the modem, as described below.

FFFFFFFF remotely reinitializes the MAC, which resumes the normal operation of the modem.

The CM MUST autonomously disable its upstream transmitter immediately upon receipt of an UP-DIS message with UP-DIS Timeout Interval = 0, regardless of any other transaction state (refer to Section 9), or the state of its control program. The modem stops all transmissions, but continues to listen to the MAC messages sent in the downstream. Once the upstream transmitter is disabled, the CM MUST cancel T3 and T4 timers and MUST not

attempt to re-establish connectivity with the CMTS until such time as i) the CM is power cycled, ii) downstream SYNC is lost for a period of time equal to the T4 timeout, or iii) reception of an UP-DIS message with a Timeout Interval = FFFFFFFF. All other UP-DIS messages MUST be ignored when the upstream is disabled.

If supported, the CM MUST autonomously reset its upstream transmitter upon receipt of an UP-DIS message with UP-DIS Timeout Interval = FFFFFFFF, regardless of any other transaction state (refer to Section 9), or the state of its control program. Resetting allows the modem to resume transmissions.

Additional, non-zero timeout values in the UP-DIS message SHOULD be supported. If supported, the CM MUST autonomously disable its upstream transmitter immediately upon receipt of an UP-DIS message with UP-DIS Timeout Interval T > 0 for a period of T milliseconds, regardless of any other transaction state (refer to Section 9), or the state of its control program. Although the timeout T is specified in milliseconds, the CM MAY extend the specified timeout by up to 100 msec. When timeout expires, the CM SHOULD reinitialize MAC as appropriate, starting with the initial ranging process and registration, because there is no guarantee that the CMTS has not de-registered it. In the disabled state, all other UP-DIS messages MUST be ignored, except for an UP-DIS message with UP-DIS Timeout Interval = FFFFFFFF or 00000000.

7 Media Access Control Protocol Operation

7.1 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 7-1.

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.

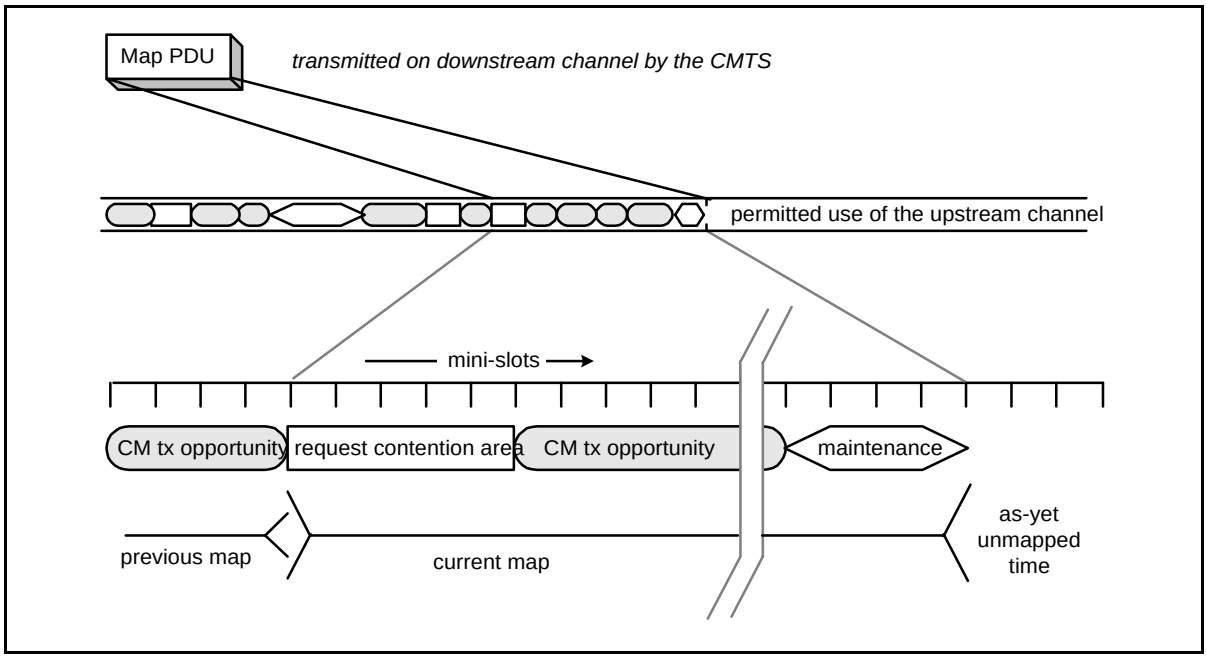


Figure 7-1. Allocation Map

The bandwidth allocation includes the following basic elements:

- Each CM has one or more short (14-bit) service identifiers (SIDs) 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.

7.1.1 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 6.3.4. 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 6.3.3. The relationship between alloc start/ack time counters and the timestamp counter is further described in Section 7.3.4.

7.1.2 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 6.3.4. 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 6-20.

Four types of Service IDs are defined:

1. 0x3FFF - broadcast, intended for all stations
2. 0x2000-0x3FFE - multicast, purpose is defined administratively. Refer to 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

All of the Information Elements defined below MUST be supported by conformant CMs. Conformant CMTSs MAY use any of these Information Elements when creating Bandwidth Allocation Maps.

7.1.2.1 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 7.4 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 Quality of Service scheduling scheme (refer to Section 8.2). Packets transmitted in this interval MUST use the Request MAC Frame format (refer to Section 6.2.5.3).

A small number of Priority Request SIDs are defined in Appendix A. These allow contention for Request IEs to be limited to service flows of a given Traffic Priority (refer to Appendix C.2.2.5.1).

7.1.2.2 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 transmit opportunities 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 6-13).

7.1.2.3 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 7.3.3), 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 (refer to Section 6.3.5).

7.1.2.4 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 6.3.5).

7.1.2.5 Short and Long Data Grant IEs

The Short and Long Data Grant IEs provide an opportunity for a CM to transmit one or more upstream PDUs. These IEs are 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). 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).

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 burst profiles are defined in the UCD, then 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 data grants pending and data acknowledgments.

7.1.2.6 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 data grants pending and data acknowledgments.

7.1.2.7 Expansion IE

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

7.1.2.8 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.

7.1.3 Requests

Requests refer to the mechanism that CMs use to indicate to the CMTS that it needs upstream bandwidth allocation. A Request MAY come as a stand-alone Request Frame transmission (refer to Section 6.2.5.3 or it MAY come as a piggyback request in the EHDR of another Frame transmission (refer to Section 6.2.6).

The Request Frame MAY be transmitted during any of the following intervals:

- Request IE
- Request/Data IE
- Short Data Grant IE
- Long Data Grant IE

A piggyback request MAY be contained in the following Extended Headers:

- Request EH element
- Upstream Privacy EH element
- Upstream Privacy EH element with Fragmentation

The request MUST include:

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

The CM MUST request the number of mini-slots needed to transmit an entire frame, or a fragment containing the entire remaining portion of a frame that a previous grant has caused to be fragmented. The frame may be a single MAC frame, or a MAC frame that has been formed by the concatenation of multiple MAC frames (see Section 6.2.5.5). The request from the CM MUST be large enough to accommodate the entire necessary physical layer overhead (see Section 4.2) for transmitting the MAC frame or fragment. The CM MUST NOT make a request that would violate the limits on data grant sizes in the UCD message (see Section 6.3.3) or any limits established by QOS parameters associated with the Service Flow.

The CM MUST have only one request outstanding at a time per Service ID. If the CMTS does not immediately respond with a Data Grant, the CM is able to unambiguously determine that its request is still pending because the CMTS MUST continue to issue a Data Grant Pending in every MAP for as long as a request is unsatisfied.

In MAPs, the CMTS MUST NOT make a data grant greater than 255 mini-slots to any assigned Service ID. This puts an upper bound on the grant size the CM has to support.

7.1.4 Information Element Feature Usage Summary

The following table summarizes what types of frames the CM can transmit using each of the MAP IE types that represent transmit opportunities. A “MUST” entry in the table means that, if appropriate, a compliant CM implementation has to be able to transmit that type of frame in that type of opportunity. A “MAY” entry means that compliant CM implementation does not have to be able to transmit that type of frame in that type of opportunity but that it is legal for it to do so, if appropriate. A “MUST NOT” entry means that a compliant CM will never transmit that type of frame in that type of opportunity.

Table 7-1. IE Feature Compatibility Summary

Information Element	Transmit Request Frame	Transmit Concatenated MAC Frame	Transmit Fragmented MAC Frame	Transmit RNG-REQ	Transmit Any other MAC Frame
Request IE	MUST	MUST NOT	MUST NOT	MUST NOT	MUST NOT
Request/Data IE	MUST	MAY	MUST NOT	MUST NOT	MAY
Initial Maintenance IE	MUST NOT	MUST NOT	MUST NOT	MUST	MUST NOT
Station Maintenance IE	MUST NOT	MUST NOT	MUST NOT	MUST	MUST NOT
Short Data Grant IE	MAY	MUST	MUST	MUST NOT	MUST
Long Data Grant IE	MAY	MUST	MUST	MUST NOT	MUST

7.1.5 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 micro-seconds.
- Queuing delays within the CMTS — implementation-specific.
- Processing delays within the CMs — MUST allow a minimum processing time by each CM as specified in Appendix 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. A CM MUST be able to support multiple outstanding MAPs. Even though multiple 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.

7.1.6 Protocol Example

This section illustrates the interchange between the CM and the CMTS when the CM has data to transmit (Figure 7-2). Suppose a given CM has a data PDU available for transmission.

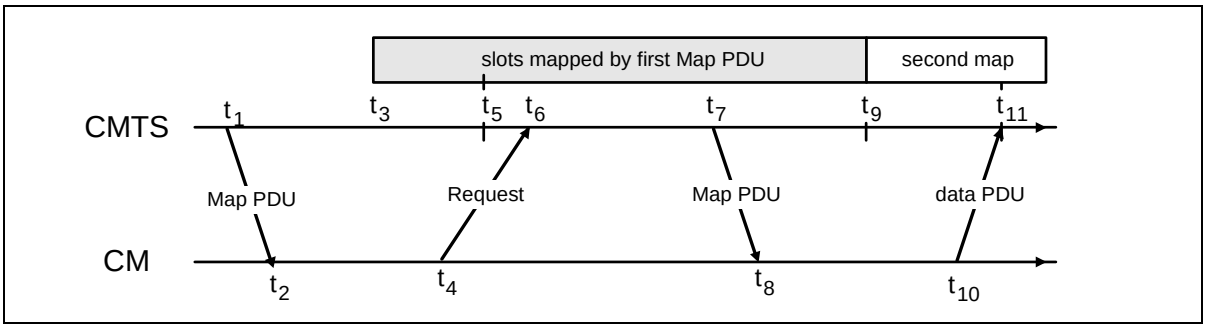


Figure 7-2. Protocol Example

Description

1. 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).
2. 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 7.4, also the multicast SID definitions in Section A.2).
3. 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 7.3.3) so that the request will arrive at the CMTS at t_6.
4. 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.)
5. 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}.
6. At t_8, the CM receives the MAP and scans for its data grant.

7. 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. (Refer to Section 7.4.1)

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.

7.2 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, then the CMTS MUST send one allocation MAP per upstream channel. The MAP's channel identifier, taken with the Upstream Channel Descriptor Message (see Section 6.3.3), MUST specify to which channel each MAP applies. There is no requirement that the maps be synchronized across channels. Appendix H 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, then the MAP MUST be transmitted on that channel. This may necessitate that multiple copies of the same MAP be transmitted. The Alloc Start Time 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.

7.3 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.

7.3.1 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 7.3.3), the model assumes that the timing delays through the remainder of the PHY layer MUST be relatively constant with the exception of the timing offsets specified in Section 6.3.7 related to symbol rate changes to accommodate a legacy DOCSIS upstream receiver implementation. 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.

7.3.2 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 9.2.2). 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 (refer to Appendix B) has elapsed without a valid SYNC message, a cable modem MUST NOT use the upstream and MUST try to re-establish synchronization again.

7.3.3 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 with the exception of the timing offsets specified in Section 6.3.7 related to symbol rate changes to accommodate a legacy DOCSIS upstream receiver implementation. 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 7.1.2.4. 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 or the CMTS aborts ranging. Once successfully ranged, the cable modem MUST join normal data traffic in the upstream. See Section 9 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 9.2.4.

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.

7.3.4 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 7-2 relates mini-slots to the SYNC time ticks:

Table 7-2. 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

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: 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.

7.4 Upstream Transmission and 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.

This section provides an overview of upstream transmission and contention resolution. For simplicity, it refers to the decisions a CM makes, however, this is just a pedagogical tool. Since a CM can have multiple upstream Service Flows (each with its own SID) it makes these decisions on a per service queue or per SID basis. Refer to Appendix K for a state transition diagram and more detail.

7.4.1 Contention Resolution Overview

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.

Every time a CM wants to transmit in a contention region, it MUST enter the contention resolution process by setting its internal backoff window equal to the Data Backoff Start defined in the MAP currently in effect.[1]

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 Data Acknowledge 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 Data Acknowledge for it and with an Ack time more recent than the time of transmission.[2] 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 Data Backoff Start = 0 and End = 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 CM transmitting a RNG-REQ in the Initial Maintenance IE MUST perform truncated binary exponential backoff using the Ranging Backoff Start and Ranging Backoff End to control the backoff window. The algorithm works similarly to data transmissions, except of the calculation of transmit opportunities which is described in the next section.

---

1. The MAP currently in effect is the MAP whose allocation start time has occurred but which includes IEs that have not occurred.

2. Data Acknowledge IEs are intended for collision detection only and is not designed for providing reliable transport (that is the responsibility of higher layers). If a MAP is lost or damaged, a CM waiting for a Data Acknowledge MUST assume that its contention data transmission was successful and MUST NOT retransmit the data packet. This prevents the CM from sending duplicate packets unnecessarily.

7.4.2 Transmit Opportunities

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 then 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, then 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 an SID of 0x3FF4 (refer to Appendix A), then 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 IEs, Short/Long Data Grant IEs and unicast Request IEs are unicast and thus are not typically associated with contention Transmit Opportunities. They represent a single dedicated, or reservation based, Transmit Opportunity.

In summary:

Table 7-3. 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.

7.4.3 CM Bandwidth Utilization

The following rules govern the response a CM makes when processing maps.

Note: These standard behaviors can be overridden by the CM's Request/Transmission Policy (refer to Appendix C.2.2.6.3):

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. A CM MUST NOT have more than one Request 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.

7.5 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.

7.5.1 MAC Messages

MAC Management Messages (Section 6.3) MUST NOT be encrypted, except for certain cases where such a frame is included in a fragmented concatenated burst on the upstream. (Refer to Section 6.2.7.1)

7.5.2 Framing

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

- 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 are carried as Data PDUs to the Cable MAC transparently.

This page intentionally left blank.

8 Quality of Service & Fragmentation

This specification introduces several new Quality of Service (QoS) related concepts not present in [DOCSIS8]. These include:

- Packet Classification & Flow Identification
- Service Flow QoS Scheduling
- Dynamic Service Establishment
- Fragmentation
- Two-Phase Activation Model

8.1 Theory of Operation

The various DOCSIS protocol mechanisms described in this document can be used to support Quality of Service (QoS) for both upstream and downstream traffic through the CM and the CMTS. This section provides an overview of the QoS protocol mechanisms and their part in providing end-to-end QoS.

The requirements for Quality of Service include:

- A configuration and registration function for pre-configuring CM-based QoS **Service Flows** and traffic parameters.
- A signaling function for dynamically establishing QoS-enabled Service Flows and traffic parameters
- A traffic-shaping and traffic-policing function for Service Flow-based traffic management, performed on traffic arriving from the upper layer service interface and outbound to the RF.
- Utilization of MAC scheduling and traffic parameters for upstream Service Flows.
- Utilization of QoS traffic parameters for downstream Service Flows.
- Classification of packets arriving from the upper layer service interface to a specific active Service Flow.
- Grouping of Service Flow properties into named **Service Classes**, so upper layer entities and external applications (at both the CM and CMTS) can request Service Flows with desired QoS parameters in a globally consistent way.

The principal mechanism for providing enhanced QoS is to classify packets traversing the RF MAC interface into a **Service Flow**. A Service Flow is a unidirectional flow of packets that is provided a particular Quality of Service. The CM and CMTS provide this QoS by shaping, policing, and prioritizing traffic according to the **QoS Parameter Set** defined for the Service Flow.

The primary purpose of the Quality of Service features defined here is to define transmission ordering and scheduling on the Radio Frequency Interface. However, these features often need to work in conjunction with mechanisms beyond the RF interface in order to provide end-to-end QoS or to police the behavior of cable modems. For example, the following behaviors are permitted:

- Policies may be defined by CM MIBs which overwrite the TOS byte. Such policies are outside the scope of the RFI specification. In the upstream direction the CMTS polices the TOS byte setting regardless of how the TOS byte is derived or by whom it is written (originator or CM policy).
- The queueing of Service Flow packets at the CMTS in the downstream direction may be based on the TOS byte.
- Downstream Service Flows can be reclassified by the CM to provide enhanced service onto the subscriber-side network.

Service Flows exist in both the upstream and downstream direction, and may exist without actually being activated to carry traffic. Service Flows have a 32-bit **Service Flow Identifier** (SFID) assigned by the CMTS. All Service Flows have an SFID; active and admitted upstream Service Flows also have a 14-bit Service Identifier (SID).

At least two Service Flows must be defined in each configuration file: one for upstream and one for downstream service. The first upstream Service Flow describes the **Primary Upstream Service Flow**, and is the default Service Flow used for otherwise unclassified traffic, including both MAC Management messages and Data PDUs. The first downstream Service Flow describes service to the Primary Downstream Service Flow. Additional Service Flows defined in the Configuration file create Service Flows that are provided QoS services.

Conceptually, incoming packets are matched to a **Classifier** that determines to which QoS Service Flow the packet is forwarded. The Classifier can examine the LLC header of the packet, the IP/TCP/UDP header of the packet or some combination of the two. If the packet matches one of the Classifiers, it is forwarded to the Service Flow indicated by the SFID attribute of the Classifier. If the packet is not matched to a Classifier, it is forwarded on the Primary Service Flow.

8.1.1 Concepts

8.1.1.1 Service Flows

A **Service Flow** is a MAC-layer transport service that provides unidirectional transport of packets either to upstream packets transmitted by the CM or to downstream packets transmitted by the CMTS[1]. A Service Flow is characterized by a set of **QoS Parameters** such as latency, jitter, and throughput assurances. In order to standardize operation between the CM and CMTS, these attributes include details of how the CM requests upstream minislots and the expected behavior of the CMTS upstream scheduler.

A Service Flow is partially characterized by the following attributes[2]:

- **ServiceFlowID**: exists for all service flows
- **ServiceID**: only exists for admitted or active upstream service flows
- **ProvisionedQosParamSet**: defines a set of QoS Parameters which appears in the configuration file and is presented during registration. This MAY define the initial limit for authorizations allowed by the authorization module. The ProvisionedQosParamSet is defined once when the Service Flow is created via registration.[3]
- **AdmittedQosParamSet**: defines a set of QoS parameters for which the CMTS (and possibly the CM) are reserving resources. The principal resource to be reserved is bandwidth, but this also includes any other memory or time-based resource required to subsequently activate the flow.
- **ActiveQosParamSet**: defines set of QoS parameters defining the service actually being provided to the Service Flow. Only an Active Service Flow may forward packets.

---

1. A Service Flow, as defined here, has no direct relationship to the concept of a “flow” as defined by the IETF’s Integrated Services (intserv) Working Group [RFC-2212]. An intserv flow is a collection of packets sharing transport-layer endpoints. Multiple intserv flows can be served by a single Service Flow. However, the Classifiers for a Service Flow may be based on 802.1P/Q criteria, and so may not involve intserv flows at all.

2. Some attributes are derived from the above attribute list. The Service Class Name is an attribute of the ProvisionedQoSParamSet. The activation state of the Service Flow is determined by the ActiveQoSParamSet. If the ActiveQoSParamSet is null then the service flow is inactive.

3. The ProvisionedQoSParamSet is null when a flow is created dynamically.

A Service Flow exists when the CMTS assigns a Service Flow ID (SFID) to it. The SFID serves as the principal identifier in the CM and CMTS for the Service Flow. A Service Flow which exists has at least an SFID, and an associated Direction.

The **Authorization Module** is a logical function within the CMTS that approves or denies every change to QoS Parameters and Classifiers associated with a Service Flow. As such it defines an "envelope" that limits the possible values of the AdmittedQoSParameterSet and ActiveQoSParameterSet.

The relationship between the QoS Parameter Sets is as shown in Figure 8-1 and Figure 8-2. The ActiveQoSParameterSet is always a subset[4] of the AdmittedQoSParameterSet which is always a subset of the authorized "envelope." In the dynamic authorization model, this envelope is determined by the Authorization Module (labeled as the AuthorizedQoSParameterSet). In the provisioned authorization model, this envelope is determined by the ProvisionedQoSParameterSet. (Refer to Section 8.1.4 for further information on the authorization models)

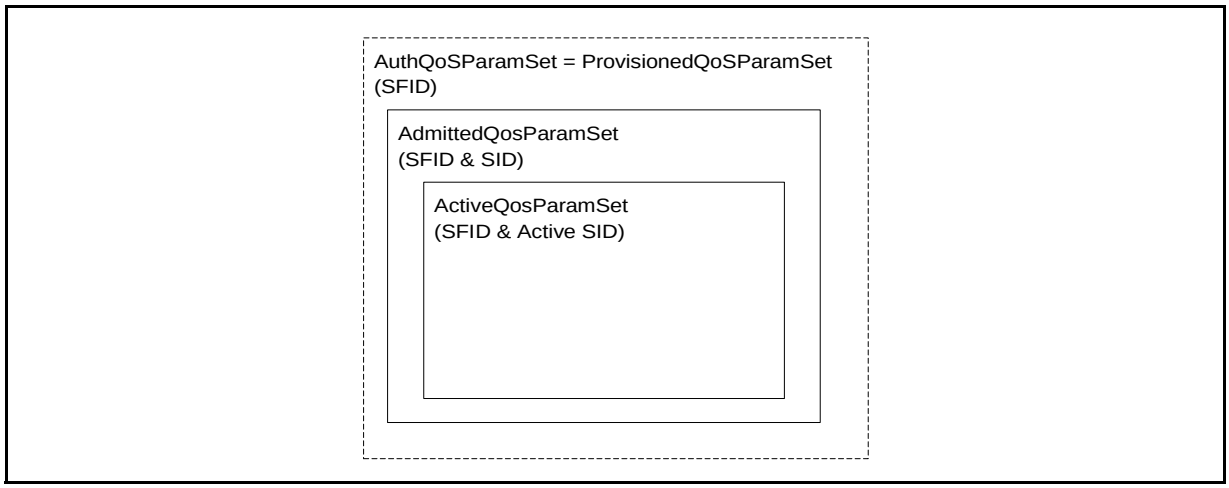


Figure 8-1. Provisioned Authorization Model "Envelopes"

4. To say that QoS Parameter Set A is a subset of QoS Parameter Set B the following MUST be true for all QoS Parameters in A and B:

if (a smaller QoS parameter value indicates less resources, *e.g.*, Maximum Traffic Rate)
A is a subset of B if the parameter in A less than or equal to the same parameter in B
if (a larger QoS parameter value indicates less resources, *e.g.*, Tolerated Grant Jitter)
A is a subset of B if the parameter in A is greater than or equal to the same parameter in B
if (the QoS parameter specifies a periodic interval, *e.g.*, Nominal Grant Interval),
A is a subset of B if the parameter in A is an integer multiple of the same parameter in B
if (the QoS parameter is not quantitative, *e.g.*, Service Flow Scheduling Type)
A is a subset of B if the parameter in A is equal to the same parameter in B

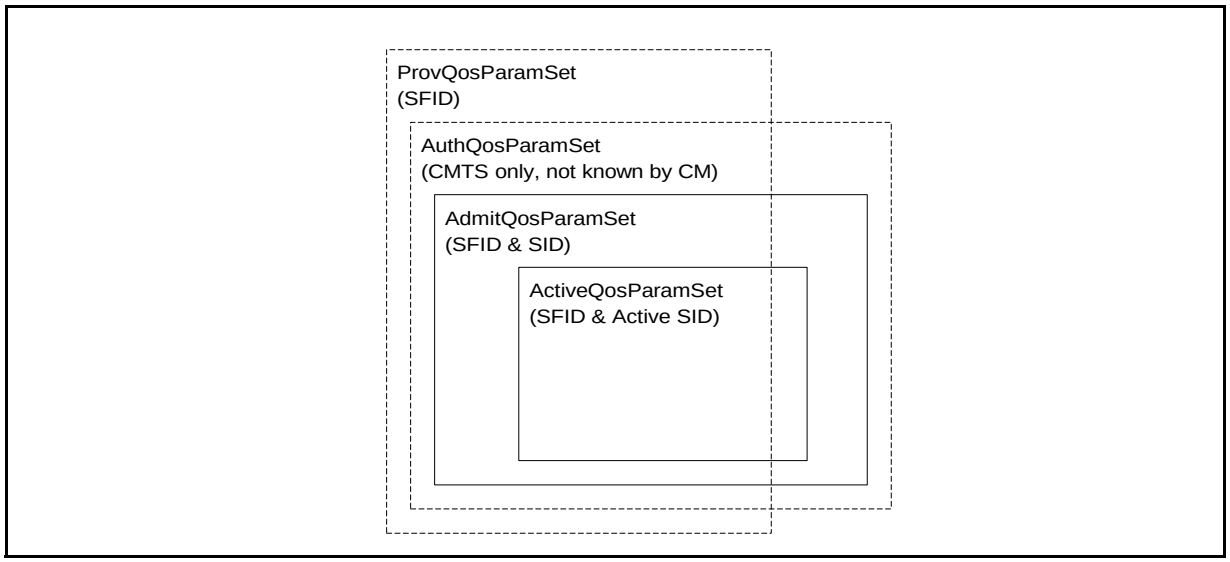


Figure 8-2. Dynamic Authorization Model "Envelopes"

It is useful to think of three types of Service Flows:

- **Provisioned:** this type of Service Flow is known via provisioning through the configuration file, its AdmittedQoSParamSet and ActiveQoSParamSet are both null. A **Provisioned Service Flow** may or may not have associated Classifiers. If a Provisioned Service Flow has associated Classifiers, the Classifiers MUST NOT be used to classify packets onto the flow, regardless of the Classifier's Activation State.
- **Admitted:** this type of Service Flow has resources reserved by the CMTS for its AdmittedQoSParamSet, but these parameters are not active (its ActiveQoSParamSet is null). **Admitted Service Flows** may have been provisioned or may have been signalled by some other mechanism. Generally, Admitted Service Flows have associated Classifiers, however, it is possible for Admitted Service Flows to use policy-based classification. If Admitted Service Flows have associated Classifiers, the classifiers MUST NOT be used to classify packets onto the flow, regardless of the classifier's activation state.
- **Active:** this type of Service Flow has resources committed by the CMTS for its QoS Parameter Set, (*e.g.*, is actively sending MAPs containing unsolicited grants for a UGS-based service flow). Its ActiveQoSParamSet is non-null. Generally, Active Service Flows have associated Classifiers, however, it is possible for Active Service Flows to use policy-based classification. Primary Service Flows may have associated Classifiers(s), but in addition to any packets matching such Classifiers, all packets that fail to match any Classifier will be sent on the Primary Service Flow for that direction.

8.1.1.2 Classifiers

A **Classifier** is a set of matching criteria applied to each packet entering the cable network. It consists of some packet matching criteria (destination IP address, for example), a **classifier priority**, and a reference to a service flow. If a packet matches the specified packet matching criteria, it is then delivered on the referenced service flow.

Several Classifiers may all refer to the same Service Flow. The classifier priority is used for ordering the application of Classifiers to packets. Explicit ordering is necessary because the patterns used by Classifiers may overlap. The priority need not be unique, but care must be taken within a classifier priority to prevent ambiguity in classification. (Refer to Section 8.1.6.1) **Downstream Classifiers** are applied by the CMTS to packets it is transmitting, and **Upstream Classifiers** are applied at the CM and may be applied at the CMTS to police the classification of upstream packets. Figure 8-3 illustrates the mappings discussed above.

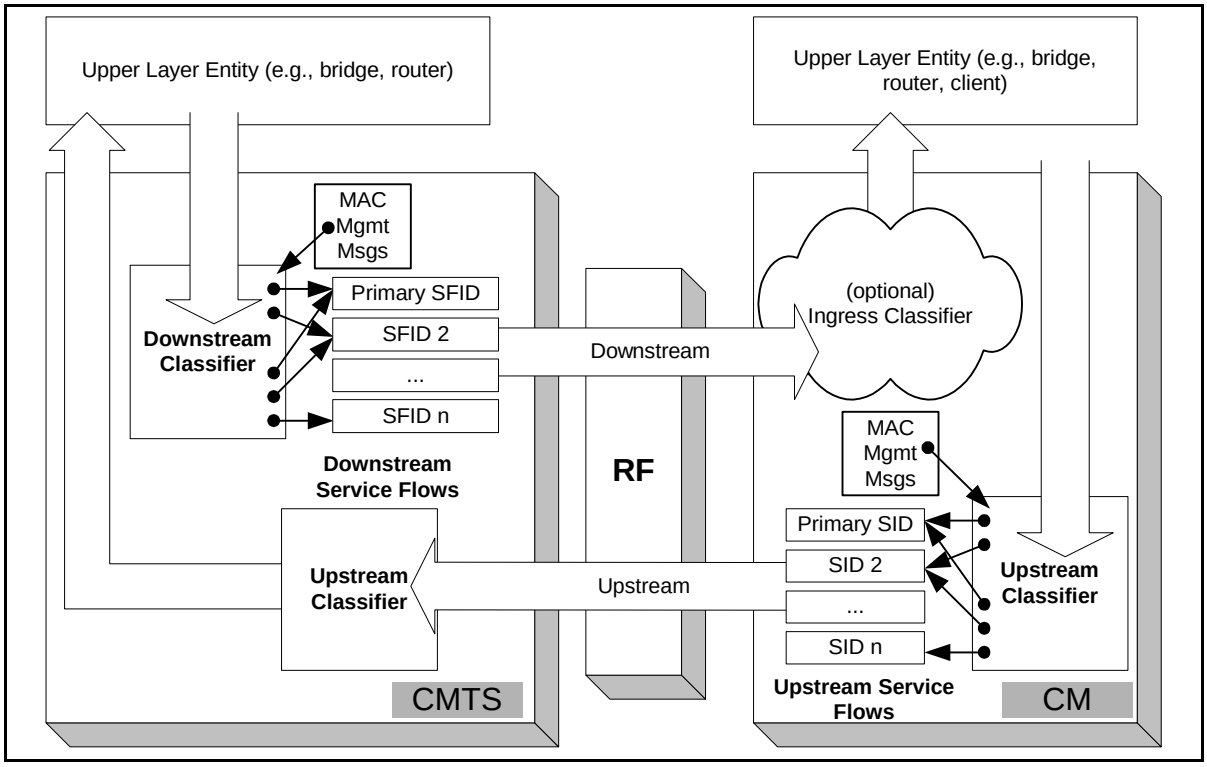


Figure 8-3. Classification within the MAC Layer

CM and CMTS Packet Classification consists of multiple Classifiers. Each Classifier contains a priority field which determines the search order for the Classifier. The highest priority Classifier MUST be applied first. If a Classifier is found that has at least one relevant parameter and all relevant parameters match the packet, the Classifier MUST forward the packet to the corresponding Service Flow (irrelevant parameters - as defined in Appendix C.2.1 - have no impact on packet classification decisions). If a Classifier contains no relevant parameters for a given packet (i.e., all parameters are irrelevant), then that packet cannot match the Classifier, and the Classifier MUST NOT forward the packet to the corresponding Service Flow. If a packet does not match any Classifier and as a result has not been classified to any other flow, then it MUST be classified to the Primary Service Flow.

The packet classification table contains the following fields:

- Priority — determines the search order for the table. Higher priority Classifiers are searched before lower priority Classifiers.

- IP Classification Parameters — zero or more of the IP classification parameters (IP TOS Range/Mask, IP Protocol, IP Source Address/Mask, IP Destination Address/Mask, TCP/UDP Source Port Start, TCP/UDP Source Port End, TCP/UDP Destination Port Start, TCP/UCP Destination Port End)
- LLC Classification Parameters — zero or more of the LLC classification parameters (Destination MAC Address, Source MAC Address, Ethertype/SAP)
- IEEE 802.1P/Q Parameters — zero or more of the IEEE classification parameters (802.1P Priority Range, 802.1Q VLAN ID)
- Service Flow Identifier — identifier of a specific flow to which this packet is to be directed

Classifiers can be added to the table either via management operations (configuration file, registration) or via dynamic operations (dynamic signaling, DOCSIS MAC sublayer service interface). SNMP-based operations can view Classifiers that are added via dynamic operations, but can not modify or delete Classifiers that are created by dynamic operations. The format for classification table parameters defined in the configuration file, registration message, or dynamic signaling message is contained in Appendix C.

Classifier attributes include an activation state (see Appendix C.2.1.3.6). The 'inactive' setting may be used to reserve resources for a classifier which is to be activated later. The actual activation of the classifier depends both on this attribute and on the state of its service flow. If the service flow is not active then the classifier is not used, regardless of the setting of this attribute.

8.1.2 Object Model

The major objects of the architecture are represented by named rectangles in Figure 8-4. Each object has a number of attributes; the attribute names which uniquely identify the object are underlined. Optional attributes are denoted with brackets. The relationship between the number of objects is marked at each end of the association line between the objects. For example, a Service Flow may be associated with from 0 to 65535 Classifiers, but a Classifier is associated with exactly one Service flow.

The Service Flow is the central concept of the MAC protocol. It is uniquely identified by a 32-bit Service Flow ID (SFID) assigned by the CMTS. Service Flows may be in either the upstream or downstream direction. A unicast Service Identifier (SID) is a 14 bit index, assigned by the CMTS, which is associated with one and only one Admitted Upstream Service Flow.

Typically, an outgoing user data Packet is submitted by an upper layer protocol (such as the forwarding bridge of a CM) for transmission on the Cable MAC interface. The packet is compared against a set of Classifiers. The matching Classifier for the Packet identifies the corresponding Service Flow via the Service Flow ID (SFID). In the case where more than one Classifier matches the packet, the highest Priority Classifier is chosen.

The Classifier matching a packet may be associated with a Payload Header Suppression Rule. A PHS Rule provides details on how header bytes of a Packet PDU can be omitted, replaced with a Payload Header Suppression Index for transmission and subsequently regenerated at the receiving end. PHS Rules are indexed by the combination of {SFID, PHSI} (refer to Section 8.4). When a Service Flow is deleted, all Classifiers and any associated PHS Rules referencing it MUST also be deleted.

The Service Class is an optional object that MAY be implemented at the CMTS. It is referenced by an ASCII name which is intended for provisioning purposes. A Service Class is defined in the CMTS to have a particular QoS Parameter Set. A Service Flow may contain a reference to the Service Class Name that selects all of the QoS parameters of the Service Class. The Service Flow QoS Parameter Sets may augment and even override the QoS parameter settings of the Service Class, subject to authorization by the CMTS. (Refer to Appendix C.2.2.5)

If a Packet has already been determined by upper layer policy mechanisms to be associated with a particular Service Class Name/Priority combination, that combination associates the packet with a particular Service Flow

directly (refer to Section 8.1.6.1). The upper layer may also be aware of the particular Service Flows in the MAC Sublayer, and may have assigned the Packet directly to a Service Flow. In these cases, a user data Packet is considered to be directly associated with a Service Flow as selected by the upper layer. This is depicted with the dashed arrow in Figure 8-4. (Refer to Appendix E)

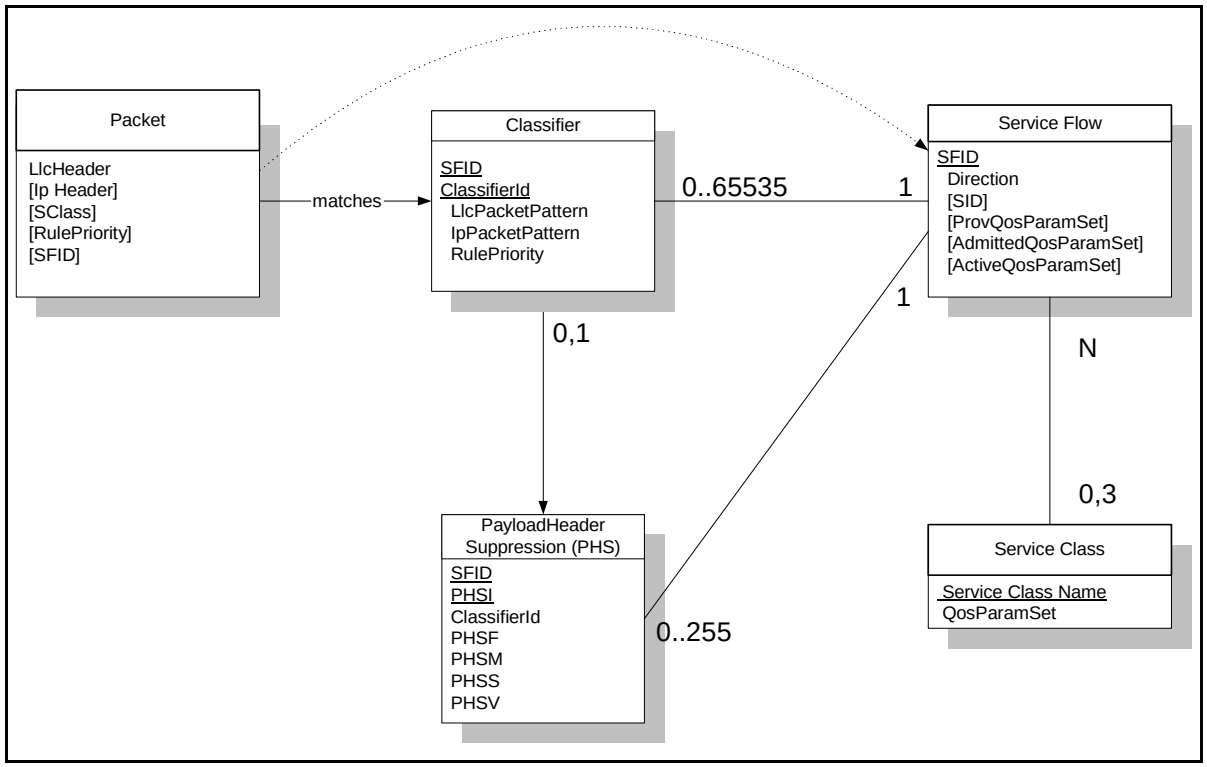


Figure 8-4. Theory of Operation Object Model

8.1.3 Service Classes

The QoS attributes of a Service Flow may be specified in two ways: either by explicitly defining all attributes, or implicitly by specifying a **Service Class Name**. A **Service Class Name** is a string which the CMTS associates with a QoS Parameter Set. It is described further below.

The Service Class serves the following purposes:

1. It allows operators, who so wish, to move the burden of configuring service flows from the provisioning server to the CMTS. Operators provision the modems with the Service Class Name; the implementation of the name is configured at the CMTS. This allows operators to modify the implementation of a given service to local circumstances without changing modem provisioning. For example, some scheduling parameters may need to be tweaked differently for two different CMTSs to provide the same service. As another example, service profiles could be changed by time of day.
2. It allows CMTS vendors to provide class-based-queuing if they choose, where service flows compete within their class and classes compete with each other for bandwidth.
3. It allows higher-layer protocols to create a Service Flow by its Service Class Name. For example, telephony signaling may direct the CM to instantiate any available Provisioned Service Flow of class “G711”.

4. It allows packet classification policies to be defined which refer to a desired service class, without having to refer to a particular service flow instance of that class.

Note: The Service Class is optional: the flow scheduling specification may always be provided in full; a service flow may belong to no service class whatsoever. CMTS implementations MAY treat such "unclassed" flows differently from "classed" flows with equivalent parameters.

Any Service Flow MAY have each of its QoS Parameter Sets specified in any of three ways:

- By explicitly including all traffic parameters
- By indirectly referring to a set of traffic parameters by specifying a Service Class Name
- By specifying a Service Class Name along with modifying parameters

The Service Class Name is "expanded" to its defined set of parameters at the time the CMTS successfully admits the Service Flow. The Service Class expansion can be contained in the following CMTS-originated messages: Registration Response, DSA-REQ, DSC-REQ, DSA-RSP and DSC-RSP. In all of these cases, the CMTS MUST include a Service Flow Encoding that includes the Service Class Name and the QoS Parameter Set of the Service Class. If a CM-initiated request contained any supplemental or overriding Service Flow parameters, a successful response MUST also include these parameters.

When a Service Class name is given in an admission or activation request, the returned QoS Parameter Set may change from activation to activation. This can happen because of administrative changes to the Service Class' QoS Parameter Set at the CMTS. If the definition of a Service Class Name is changed at the CMTS (*e.g.*, its associated QoS Parameter Set is modified), it has no effect on the QoS Parameters of existing Service Flows associated with that Service Class. A CMTS MAY initiate DSC transactions to existing Service Flows which reference the Service Class Name to affect the changed Service Class definition.

When a CM uses the Service Class Name to specify the Admitted QoS Parameter Set, the expanded set of TLV encodings of the Service Flow will be returned to the CM in the response message (REG-RSP, DSA-RSP, or DSC-RSP). Use of the Service Class Name later in the activation request may fail if the definition of the Service Class Name has changed and the new required resources are not available. Thus, the CM SHOULD explicitly request the expanded set of TLVs from the response message in its later activation request.

8.1.4 Authorization

Every change to the Service Flow QoS Parameters MUST be approved by an authorization module. This includes every REG-REQ or DSA-REQ message to create a new Service Flow, and every DSC-REQ message to change a QoS Parameter Set of an existing Service Flow. Such changes include requesting an admission control decision (*e.g.*, setting the AdmittedQoSParamSet) and requesting activation of a Service Flow (*e.g.*, setting the ActiveQoSParameterSet). Reduction requests regarding the resources to be admitted or activated are also checked by the authorization module, as are requests to add or change the Classifiers.

In the static authorization model, the authorization module receives all registration messages, and stores the provisioned status of all "deferred" Service Flows. Admission and activation requests for these provisioned service flows will be permitted, as long as the Admitted QoS Parameter Set is a subset of the Provisioned QoS Parameter Set, and the Active QoS Parameter Set is a subset of the Admitted QoS Parameter Set. Requests to change the Provisioned QoS Parameter Set will be refused, as will requests to create new dynamic Service Flows. This defines a static system where all possible services are defined in the initial configuration of each CM.

In the dynamic authorization model, the authorization module not only receives all registration messages, but also communicates through a separate interface to an independent policy server. This policy server may provide to the authorization module advance notice of upcoming admission and activation requests, and specifies the proper authorization action to be taken on those requests. Admission and activation requests from a CM are then

checked by the Authorization Module to ensure that the ActiveQoSParameterSet being requested is a subset of the set provided by the policy server. Admission and activation requests from a CM that are signalled in advance by the external policy server are permitted. Admission and activation requests from a CM that are not pre-signalled by the external policy server may result in a real-time query to the policy server, or may be refused.

During registration, the CM MUST send to the CMTS the authenticated set of TLVs derived from its configuration file which defines the Provisioned QoS Parameter Set. Upon receipt and verification at the CMTS, these are handed to the Authorization Module within the CMTS. The CMTS MUST be capable of caching the Provisioned QoS Parameter Set, and MUST be able to use this information to authorize dynamic flows which are a subset of the Provisioned QoS Parameter Set. The CMTS SHOULD implement mechanisms for overriding this automated approval process (such as described in the dynamic authorization model). For example:

- Deny all requests whether or not they have been pre-provisioned
- Define an internal table with a richer policy mechanism but seeded by the configuration file information
- Refer all requests to an external policy server

8.1.5 Types of Service Flows

It is useful to think about three basic types of Service Flows. This section describes these three types of Service Flows in more detail. However, it is important to note that there are more than just these three basic types. (Refer to Appendix C.2.2.3.5)

8.1.5.1 Provisioned Service Flows

A Service Flow may be Provisioned but not immediately activated (sometimes called “deferred”). That is, the description of any such service flow in the TFTP configuration file contains an attribute which provisions but defers activation and admission (refer to Appendix C.2.2.3.5). During Registration, the CMTS assigns a Service Flow ID for such a service flow but does not reserve resources. The CMTS MAY also require an exchange with a policy module prior to admission.

As a result of external action beyond the scope of this specification (*e.g.*, [PKTCBL-MGCP]), the CM MAY choose to activate a Provisioned Service Flow by passing the Service Flow ID and the associated QoS Parameter Sets. The CM MUST also provide any applicable Classifiers. If authorized and resources are available, the CMTS MUST respond by assigning a unique unicast SID for the upstream Service Flow. The CMTS MAY deactivate the Service Flow, but SHOULD NOT delete the Service Flow during the CM registration epoch.

As a result of external action beyond the scope of this specification (*e.g.*, [PKTCBL-MGCP]), the CMTS MAY choose to activate a Service Flow by passing the Service Flow ID as well as the SID and the associated QoS Parameter Sets. The CMTS MUST also provide any applicable Classifiers. The CMTS MAY deactivate the Service Flow, but SHOULD NOT delete the Service Flow during the CM registration epoch. Such a Provisioned Service Flow MAY be activated and deactivated many times (through DSC exchanges). In all cases, the original Service Flow ID MUST be used when reactivating the service flow.

8.1.5.2 Admitted Service Flows

This protocol supports a two-phase activation model which is often utilized in telephony applications. In the two-phase activation model, the resources for a "call" are first "admitted," and then once the end-to-end negotiation is completed (*e.g.*, called party's gateway generates an "off-hook" event) the resources are "activated." Such a two-phase model serves the purposes of a) conserving network resources until a complete end-to-end connection has been established, b) performing policy checks and admission control on resources as quickly as possible, and, in particular, before informing the far end of a connection request, and c) preventing several potential theft-of-service scenarios.

For example, if an upper layer service were using unsolicited grant service, and the addition of upper-layer flows could be adequately provided by increasing the Grants Per Interval QoS parameter, then the following might be used. When the first upper-layer flow is pending, the CM issues a DSA-Request with the Admit Grants Per Interval parameter equal one, and the Activate Grants Per Interval parameter equal zero. Later when the upper-layer flow becomes active, it issues a DSC-Request with the instance of the Activate Grants-per-Interval parameter equal to one. Admission control was performed at the time of the reservation, so the later DSC-Request, having the Activate parameters within the range of the previous reservation, is guaranteed to succeed. Subsequent upper-layer flows would be handled in the same way. If there were three upper-layer flows establishing connections, with one flow already active, the Service Flow would have Admit(ted) Grants-per-Interval equal four, and Active Grants-per-Interval equal one.

An activation request of a Service Flow where the new ActiveQoSParamSet is a subset of the AdmittedQoS-ParamSet and no new classifiers are being added MUST be allowed (except in the case of catastrophic failure). An admission request where the AdmittedQoSParamSet is a subset of the previous AdmittedQoSParamSet, so long as the ActiveQoSParamSet remains a subset of the AdmittedQoSParameterSet, MUST succeed.

A Service Flow that has resources assigned to its AdmittedQoSParamSet, but whose resources are not yet completely activated, is in a transient state. A timeout value MUST be enforced by the CMTS that requires Service Flow activation within this period. (Refer to Appendix C.2.2.5.7) If Service Flow activation is not completed within this interval, the assigned resources in excess of the active QoS parameters MUST be released by the CMTS.

It is possible in some applications that a long-term reservation of resources is necessary or desirable. For example, placing a telephone call on hold should allow any resources in use for the call to be temporarily allocated to other purposes, but these resources must be available for resumption of the call later. The AdmittedQoSParamSet is maintained as "soft state" in the CMTS; this state must be refreshed periodically for it to be maintained without the above timeout releasing the non-activated resources. This refresh MAY be signalled with a periodic DSC-REQ message with identical QoS Parameter Sets, or MAY be signalled by some internal mechanism within the CMTS outside of the scope of this specification (*e.g.*, by the CMTS monitoring RSVP refresh messages). Every time a refresh is signaled to the CMTS, the CMTS MUST refresh the "soft state."

8.1.5.3 Active Service Flows

A Service Flow that has a non-NULL set of ActiveQoSParameters is said to be an Active Service Flow. It is requesting[5] and being granted bandwidth for transport of data packets. An admitted Service Flow may be made active by providing an ActiveQoSParameterSet, signaling the resources actually desired at the current time. This completes the second stage of the two-phase activation model. (Refer to Section 8.1.5.2)

A Service Flow may be Provisioned and immediately activated. This is the case for the Primary Service Flows. It is also typical of Service Flows for monthly subscription services, etc. These Service Flows are established at registration time and MUST be authorized by the CMTS based on the CMTS MIC. These Service Flows MAY also be authorized by the CMTS authorization module.

Alternatively, a Service Flow may be created dynamically and immediately activated. In this case, two-phase activation is skipped and the Service Flow is available for immediate use upon authorization.

8.1.6 Service Flows and Classifiers

The basic model is that the Classifiers associate packets into exactly one Service Flow. The Service Flow Encodings provide the QoS Parameters for treatment of those packets on the RF interface. These encodings are described in Appendix C.2.

In the upstream direction, the CM MUST classify upstream packets to Active Service Flows. The CMTS MUST classify downstream traffic to Active Downstream Service Flows. There MUST be a default downstream service flow for otherwise unclassified broadcast and multicast traffic.

The CMTS polices packets in upstream Service Flows to ensure the integrity of the QoS Parameters and the packet's TOS value. When the rate at which packets are sent is greater than the policed rate at the CMTS, then these packets MAY be dropped by the CMTS (refer to Appendix C.2.2.5.2). When the value of the TOS byte is incorrect, the CMTS (based on policy) MUST police the stream by overwriting the TOS byte (refer to Appendix C.2.2.6.10).

It may not be possible for the CM to forward certain upstream packets on certain Service Flows. In particular, a Service Flow using unsolicited grant service with fragmentation disabled cannot be used to forward packets larger than the grant size. If a packet is classified to a Service Flow on which it cannot be transmitted, the CM MUST either transmit the packet on the Primary Service Flow or discard the packet depending on the Request/ Transmission Policy of the Service Flow to which the packet was classified.

MAC Management messages may only be matched by a classifier that contains an "Ethertype/DSAP/MacType" (Appendix C.2.1.6.3) parameter encoding and when the "type" field of the MAC Management Message Header (Section 6.3.1) matches that parameter. One exception is that the Primary SID MUST be used for station maintenance, as specified in Section 6.1.2.3, even if a classifier matches the upstream RNG-REQ message of

5. According to its Request/Transmission Policy (refer to Appendix C.2.2.6.3)

station maintenance. In the absence of any classifier matching a MAC Management message, it SHOULD be transmitted on the Primary Service Flow. Other than those MAC message types precluded from classification in Appendix C.2.1.6.3, a CM or CMTS MAY forward an otherwise unclassified MAC message on any Service Flow in an implementation-specific manner.

Although MAC Management messages are subject to classification, they are not considered part of any service flow. Transmission of MAC Management messages MUST NOT influence any QoS calculations of the Service Flow to which they are classified. Delivery of MAC Management messages is implicitly influenced by the attributes of the associated service flow.

8.1.6.1 Policy-Based Classification and Service Classes

As noted in Appendix E, there are a variety of ways in which packets may be enqueued for transmission at the MAC Service Interface. At one extreme are embedded applications that are tightly bound to a particular Payload Header Suppression Rule (refer to Section 8.4) and which forego more general classification by the MAC. At the other extreme are general transit packets of which nothing is known until they are parsed by the MAC Classification rules. Another useful category is traffic to which policies are applied by a higher-layer entity and then passed to the MAC for further classification to a particular service flow.

Policy-based classification is, in general, beyond the scope of this specification. One example might be the docsDevFilterIpPolicyTable defined in the Cable Device MIB [RFC-2669]. Such policies may tend to be longer-lived than individual service flows and MAC classifiers and so it is appropriate to layer the two mechanisms, with a well-defined interface between policies and MAC Service Flow Classification.

The interface between the two layers is the addition of two parameters at the MAC transmission request interface. The two parameters are a Service Class Name and a Rule Priority that is applied to matching the service class name. The Policy Priority is from the same number space as the Packet Classifier Priority of the packet-matching rules used by MAC classifiers. The MAC Classification algorithm is now:

```
MAC_DATA.request( PDU,
                  ServiceClassName,
                  RulePriority)

TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)
SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
IF (SearchID not NULL and Classifier.RulePriority >= MAC_DATA.RulePriority)
      TxServiceFlowID = SearchID

IF (TxServiceFlowID = NULL)
     TRANSMIT_PDU (PrimaryServiceFlowID)
ELSE
      TRANSMIT_PDU (TxServiceFlowID)
```

While Policy Priority competes with Packet Classifier Priority and its choice might in theory be problematic, it is anticipated that well-known ranges of priorities will be chosen to avoid ambiguity. In particular, dynamically-added classifiers MUST use the priority range 64-191. Classifiers created as part of registration, as well as policy-based classifiers, may use zero through 255, but SHOULD avoid the dynamic range.

Note: Classification within the MAC sublayer is intended to simply associate a packet with a service flow. If a packet is intended to be dropped it MUST be dropped by the higher-layer entity and not delivered to the MAC sublayer.

8.1.7 General Operation

8.1.7.1 Static Operation

Static configuration of Classifiers and Service Flows uses the Registration process. A provisioning server provides the CM with configuration information. The CM passes this information to the CMTS in a Registration Request. The CMTS adds information and replies with a Registration Response. The CM sends a Registration Acknowledge to complete registration.

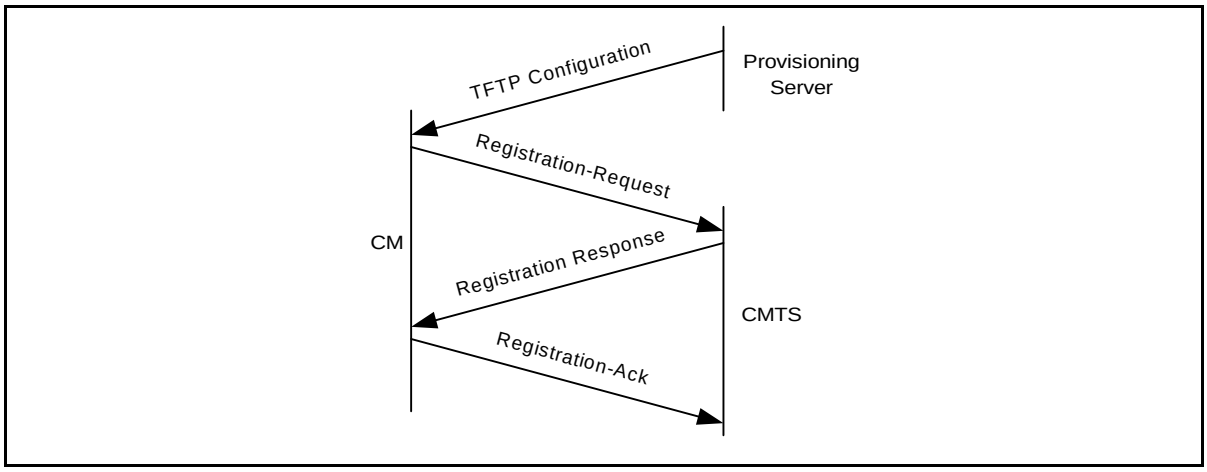


Figure 8-5. Registration Message Flow

A TFTP configuration file consists of one or more instances of Classifiers and Service Flow Encodings. Classifiers are loosely ordered by 'priority'. Each Classifier refers to a Service Flow via a 'service flow reference'. Several Classifiers may refer to the same Service Flow. Additionally, more than one Classifier may have the same priority, and in this case, the particular classifier used is not defined.

Table 8-1. TFTP File Contents

Items	Point To Service Flow Reference	Service Flow Reference	Service Flow ID
Upstream Classifiers Each containing a Service Flow Reference (pointer)	1..n		
Downstream Classifiers Each containing a Service Flow Reference (pointer)	(n+1)..q		
Service Flow Encodings Immediate activation requested, upstream		1..m	None Yet
Service Flow Encodings Provisioned for later activation requested, upstream		(m+1)..n	None Yet
Service Flow Encodings Immediate activation requested, downstream		(n+1)..p	None Yet
Service Flow Encodings Provisioned for later activation requested, downstream		(p+1)..q	None Yet

Service Flow Encodings contain either a full definition of service attributes (omitting defaultable items if desired) or a service class name. A service class name is an ASCII string which is known at the CMTS and which indirectly specifies a set of QoS Parameters. (Refer to Section 8.1.3 and Appendix C.2.2.3.4)

Note: At the time of the TFTP configuration file, Service Flow References exist as defined by the provisioning server. Service Flow Identifiers do not yet exist because the CMTS is unaware of these service flow definitions.

The Registration Request packet contains Downstream Classifiers (if to be immediately activated) and all Inactive Service Flows. The configuration file, and thus, the Registration Request, generally does not contain a Downstream Classifier if the corresponding Service Flow is requested with deferred activation. This allows for late binding of the Classifier when the Flow is activated.

Table 8-2. Registration Request Contents

Items	Point To Service Flow Reference	Service Flow Reference	Service Flow ID
Upstream Classifiers Each containing a Service Flow Reference (pointer)	1..n		
Downstream Classifiers Each containing a Service Flow Reference (pointer)	(n+1)..p		
Service Flow Encodings Immediate activation requested, upstream May specify explicit attributes or service class name		1..m	None Yet
Service Flow Encodings Provisioned for later activation requested, upstream Explicit attributes or service class name		(m+1)..n	None Yet
Service Flow Encodings Immediate activation requested, downstream Explicit attributes or service name		(n+1)..p	None Yet
Service Flow Encodings Provisioned for later activation requested, downstream Explicit attributes or service name		(p+1)..q	None Yet

The Registration Response sets the QoS Parameter Sets according to the Quality of Service Parameter Set Type in the Registration Request.

The Registration Response preserves the Service Flow Reference attribute, so that the Service Flow Reference can be associated with SFID and/or SID.

Table 8-3. Registration Response Contents

Items	Service Flow Reference	Service Flow Identifier	Service Identifier
Active Upstream Service Flows Explicit attributes	1..m	SFID	SID
Provisioned Upstream Service Flows Explicit attributes	(m+1)..n	SFID	Not Yet
Active Downstream Service Flows Explicit attributes	(n+1)..p	SFID	N/A
Provisioned Downstream Service Flows Explicit attributes	(p+1)..q	SFID	N/A

The SFID is chosen by the CMTS to identify a downstream or upstream service Flow that has been authorized but not activated. A DSC-Request from a modem to admit or activate a Provisioned Service Flow contains its SFID. If it is a downstream Flow then the Downstream Classifier is also included.

8.1.7.2 Dynamic Service Flow Creation — CM Initiated

Service Flows may be created by the Dynamic Service Addition process, as well as through the Registration process outlined above. The Dynamic Service Addition may be initiated by either the CM or the CMTS, and may create one upstream and/or one downstream dynamic Service Flow(s). A three-way handshake is used to create Service Flows. The CM-initiated protocol is illustrated in Figure 8-6 and described in detail in Section 9.4.2.1.

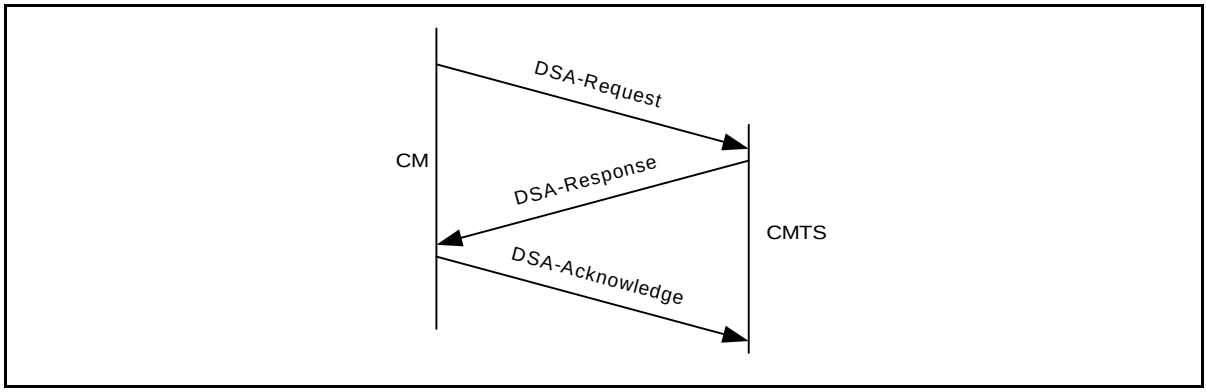


Figure 8-6. Dynamic Service Addition Message Flow — CM Initiated

A DSA-Request from a CM contains Service Flow Reference(s), QoS Parameter set(s) (marked either for admission-only or for admission and activation) and any required Classifiers.

8.1.7.3 Dynamic Service Flow Creation — CMTS Initiated

A DSA-Request from a CMTS contains Service Flow Identifier(s) for one upstream and/or one downstream Service Flow, possibly a SID, set(s) of active or admitted QoS Parameters, and any required Classifier(s). The protocol is as illustrated in Figure 8-7 and is described in detail in Section 9.4.2.2.

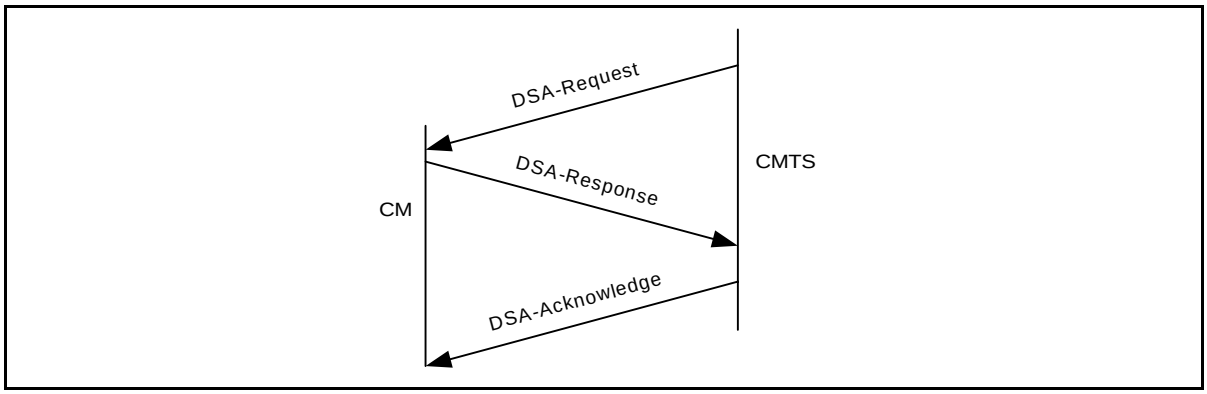


Figure 8-7. Dynamic Service Addition Message Flow — CMTS Initiated

8.1.7.4 Dynamic Service Flow Modification and Deletion

In addition to the methods presented above for creating service flows, protocols are defined for modifying and deleting service flows. Refer to Section 9.4.3 and Section 9.4.4.

Both provisioned and dynamically created Service flows are modified with the DSC message, which can change the Admitted and Active QoS Parameter sets of the flow. The DSC can also add, replace, or delete classifiers, and add, add parameters to, or delete PHS rules.

A successful DSC transaction changes a Service Flow's QoS parameters by replacing both the Admitted and Active QoS parameter sets. If the message contains only the Admitted set, the Active set is set to null and the flow is deactivated. If the message contains neither set ('000' value used for Quality of Service Parameter Set type, see Appendix C.2.2.5.1) then both sets are set to null and the flow is de-admitted. When the message contains both QoS parameter sets, the Admitted set is checked first and, if admission control succeeds, the Active set in the message is checked against the Admitted set in the message to ensure that it is a subset (see Section 8.1.1.1). If all checks are successful, the QoS parameter sets in the message become the new Admitted and Active QoS parameter sets for the Service Flow. If either of the checks fails, the DSC transaction fails and the Service Flow QoS parameter sets are unchanged.

8.2 Upstream Service Flow Scheduling Services

The following sections define the basic upstream Service Flow scheduling services and list the QoS parameters associated with each service. A detailed description of each QoS parameter is provided in Appendix C. The section also discusses how these basic services and QoS parameters can be combined to form new services, such as, Committed Information Rate (CIR) service.

Scheduling services are designed to improve the efficiency of the poll/grant process. By specifying a scheduling service and its associated QoS parameters, the CMTS can anticipate the throughput and latency needs of the upstream traffic and provide polls and/or grants at the appropriate times.

Each service is tailored to a specific type of data flow as described below. The basic services comprise: Unsolicited Grant Service (UGS), Real-Time Polling Service (rtPS), Unsolicited Grant Service with Activity Detection (UGS-AD), Non-Real-Time Polling Service (nrtPS) and Best Effort (BE) service. Table 8-4 shows the relationship between the scheduling services and the related QoS parameters.

8.2.1 Unsolicited Grant Service

The Unsolicited Grant Service (UGS) is designed to support real-time service flows that generate fixed size data packets on a periodic basis, such as Voice over IP. The service offers fixed size grants on a real-time periodic basis, which eliminate the overhead and latency of CM requests and assure that grants will be available to meet the flow's real-time needs. The CMTS MUST provide fixed size data grants at periodic intervals to the Service Flow. In order for this service to work correctly, the Request/Transmission Policy (refer to Appendix C.2.2.6.3) setting MUST be such that the CM is prohibited from using any contention request or request/data opportunities and the CMTS SHOULD NOT provide any unicast request opportunities. The Request/Transmission Policy MUST also prohibit piggyback requests. This will result in the CM only using unsolicited data grants for upstream transmission. All other bits of the Request/Transmission Policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the Unsolicited Grant Size, the Nominal Grant interval, the Tolerated Grant Jitter and the Request/Transmission Policy. (Refer to Appendix M)

The Unsolicited Grant Synchronization Header (UGSH) in the Service Flow EH Element (refer to Section 6.2.6.3.2) is used to pass status information from the CM to the CMTS regarding the state of the UGS Service Flow. The most significant bit of the UGSH is the Queue Indicator (QI) flag. When the QI flag is set it indicates a rate overrun condition for the Service Flow. When the QI flag is clear it indicates a rate non-overrun condition for the Service Flow. The QI flag allows the CMTS to provide a dynamic rate-compensation function by issuing additional grants.

The CM MUST set the QI flag when it detects that the packet reception rate is greater than the upstream transmission rate. The CM MUST clear the QI flag when it detects that the packet reception rate is equal to or less than the upstream transmission rate and the queued packet backlog is cleared.

The number of packets already queued for upstream transmission is a measure of the rate differential between received and transmitted packets. The CM SHOULD set the QI flag when the number of packets queued is greater than the number of Grants per interval parameter of the Active QoS set. The CM SHOULD clear the QI flag when the number of packets queued is less than or equal to the number of Grants per interval parameter of the Active QoS set. The QI flag of each packet MAY be set either at the time the packet is received and queued or at the time the packet is dequeued and transmitted.

The CM MAY set/clear the QI flag using a threshold of two times the number of Grants per Interval parameter of the Active QoS set. Alternatively, the CM MAY provide hysteresis by setting the QI flag using a threshold of two times the number of Grants per Interval, then clearing it using a threshold of one times the number of Grants per Interval.

The CMTS MUST NOT allocate more grants per Nominal Grant Interval than the Grants Per Interval parameter of the Active QoS Parameter Set, excluding the case when the QI bit of the UGSH is set. In this case, the CMTS SHOULD grant up to 1% additional bandwidth for clock rate mismatch compensation. If the CMTS grants additional bandwidth, it MUST limit the total number of bytes forwarded on the flow during any time interval to Max(T), as described in the expression:

Max(T) = T * (R*1.01) + 3B

Where Max(T) is the maximum number of bytes transmitted on the flow over a time T (in units of seconds), R = (grant_size * grants_per_interval)/nominal_grant_interval, and B = grant_size*grants_per_interval.

The active grants field of the UGSH is ignored with UGS service. The CMTS policing of the Service Flow remains unchanged.

8.2.2 Real-Time Polling Service

The Real-Time Polling Service (rtPS) is designed to support real-time service flows that generate variable size data packets on a periodic basis, such as MPEG video. The service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allow the CM to specify the size of the desired grant. This service requires more request overhead than UGS, but supports variable grant sizes for optimum data transport efficiency.

The CMTS MUST provide periodic unicast request opportunities. In order for this service to work correctly, the Request/Transmission Policy setting (refer to Appendix C.2.2.6.3) SHOULD be such that the CM is prohibited from using any contention request or request/data opportunities. The Request/Transmission Policy SHOULD also prohibit piggyback requests. The CMTS MAY issue unicast request opportunities as prescribed by this service even if a grant is pending. This will result in the CM using only unicast request opportunities in order to obtain upstream transmission opportunities (the CM could still use unsolicited data grants for upstream transmission as well). All other bits of the Request/Transmission Policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the Nominal Polling Interval, the Tolerated Poll Jitter and the Request/Transmission Policy.

8.2.3 Unsolicited Grant Service with Activity Detection

The Unsolicited Grant Service with Activity Detection (UGS/AD) is designed to support UGS flows that may become inactive for substantial portions of time (*i.e.*, tens of milliseconds or more), such as Voice over IP with silence suppression. The service provides Unsolicited Grants when the flow is active and unicast polls when the flow is inactive. This combines the low overhead and low latency of UGS with the efficiency of rtPS. Though USG/AD combines UGS and rtPS, only one scheduling service is active at a time.

The CMTS MUST provide periodic unicast grants, when the flow is active, but MUST revert to providing periodic unicast request opportunities when the flow is inactive. [The CMTS can detect flow inactivity by detecting unused grants. However, the algorithm for detecting a flow changing from an active to an inactive state is dependent on the CMTS implementation]. In order for this service to work correctly, the Request/Transmission Policy setting (refer to Appendix C.2.2.6.3) MUST be such that the CM is prohibited from using any contention request or request/data opportunities. The Request/Transmission Policy MUST also prohibit piggyback requests. This results in the CM using only unicast request opportunities in order to obtain upstream transmission opportunities. However, the CM will use unsolicited data grants for upstream transmission as well. All other bits of the Request/Transmission Policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the Nominal Polling Interval, the Tolerated Poll Jitter, the Nominal Grant Interval, the Tolerated Grant Jitter, the Unsolicited Grant Size and the Request/Transmission Policy.

In UGS-AD service, when restarting UGS after an interval of rtPS, the CMTS SHOULD provide additional grants in the first (and/or second) grant interval such that the CM receives a total of one grant for each grant interval from the time the CM requested restart of UGS, plus one additional grant. (Refer to Appendix M) Because the Service Flow is provisioned as a UGS flow with a specific grant interval and grant size, when restarting UGS, the CM MUST NOT request a different sized grant than the already provisioned UGS flow. As with any Service Flow, changes can only be requested with a DSC command. If the restarted activity requires more than one grant per interval, the CM MUST indicate this in the Active Grants field of the UGSH beginning with the first packet sent.

The Service Flow Extended Header Element allows for the CM to dynamically state how many grants per interval are required to support the number of flows with activity present. In UGS/AD, the CM MAY use the Queue Indicator Bit in the UGSH. The remaining seven bits of the UGSH define the Active Grants field. This field defines the number of grants within a Nominal Grant Interval that this Service Flow currently requires. When using UGS/AD, the CM MUST indicate the number of requested grants per Nominal Grant Interval in this field. The Active Grants field of the UGSH is ignored with UGS without Activity Detection. This field allows the CM to signal to the CMTS to dynamically adjust the number of grants per interval that this UGS Service Flow is actually using. The CM MUST NOT request more than the number of Grants per Interval in the ActiveQoSParameterSet.

If the CMTS allocates additional bandwidth in response to the QI bit, it MUST use the same rate limiting formula as UGS, but the formula only applies to steady state periods where the CMTS has adjusted the grants_per_interval to match the active_grants requested by the CM.

When the CM is receiving unsolicited grants and it detects no activity on the Service Flow, it MAY send one packet with the Active Grants field set to zero grants and then cease transmission. Because this packet may not be received by the CMTS, when the Service Flow goes from inactive to active the CM MUST be able to restart transmission with either polled requests or unsolicited grants.

8.2.4 Non-Real-Time Polling Service

The Non-Real-Time Polling Service (nrtPS) is designed to support non real-time service flows that require variable size data grants on a regular basis, such as high bandwidth FTP. The service offers unicast polls on a regular basis which assures that the flow receives request opportunities even during network congestion. The CMTS typically polls nrtPS SIDs on an (periodic or non-periodic) interval on the order of one second or less.

The CMTS MUST provide timely unicast request opportunities. In order for this service to work correctly, the Request/Transmission Policy setting (refer to Appendix C.2.2.6.2) SHOULD be such that the CM is allowed to use contention request opportunities. This will result in the CM using contention request opportunities as well as unicast request opportunities and unsolicited data grants. All other bits of the Request/Transmission Policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are Nominal Polling Interval, Minimum Reserved Traffic Rate, Maximum Sustained Traffic Rate, Request/Transmission Policy and Traffic Priority.

8.2.5 Best Effort Service

The intent of the Best Effort (BE) service is to provide efficient service to best effort traffic. In order for this service to work correctly, the Request/Transmission Policy setting SHOULD be such that the CM is allowed to use contention request opportunities. This will result in the CM using contention request opportunities as well as unicast request opportunities and unsolicited data grants. All other bits of the Request/Transmission Policy are not relevant to the fundamental operation of this scheduling service and should be set according to network policy. The key service parameters are the Minimum Reserved Traffic Rate, the Maximum Sustained Traffic Rate, and the Traffic Priority.

8.2.6 Other Services

8.2.6.1 Committed Information Rate (CIR)

A Committed Information Rate (CIR) service can be defined a number of different ways. For example, it could be configured by using a Best Effort service with a Minimum Reserved Traffic Rate or a nrtPS with a Minimum Reserved Traffic Rate.

8.2.7 Parameter Applicability for Upstream Service Scheduling

Table 8-4 summarizes the relationship between the scheduling services and key QoS parameters. A detailed description of each QoS parameter is provided in Appendix C.

Table 8-4. Parameter Applicability for Upstream Service Scheduling

Service Flow Parameter	Best Effort	Non-Real-Time Polling	Real-Time Polling	Unsolicited Grant	Unsolicited Grant with Activity Det.
Miscellaneous					
• Traffic Priority	Optional Default = 0	Optional Default = 0	N/A[a]	N/A	N/A
• Max Concatenated Burst	Optional	Optional	Optional	N/A	N/A
• Upstream Scheduling Service Type	Optional Default = 2	Mandatory	Mandatory	Mandatory	Mandatory
• Request/Transmission Policy	Optional Default = 0	Mandatory	Mandatory	Mandatory	Mandatory
Maximum Rate					
• Max Sustained Traffic Rate	Optional Default = 0	Optional Default = 0	Optional Default = 0	N/A	N/A
• Max Traffic Burst	Optional Dflt = 3044	Optional Dflt = 3044	Optional Dflt = 3044	N/A	N/A
Minimum Rate					
• Min Reserved Traffic Rate	Optional Default = 0	Optional Default = 0	Optional Default = 0	N/A	N/A
• Assumed Minimum … Packet Size	Optional*	Optional*	Optional*	Optional*	Optional*
Grants					
• Unsolicited Grant Size	N/A	N/A	N/A	Mandatory	Mandatory
• Grants per Interval	N/A	N/A	N/A	Mandatory	Mandatory
• Nominal Grant Interval	N/A	N/A	N/A	Mandatory	Mandatory
• Tolerated Grant Jitter	N/A	N/A	N/A	Mandatory	Mandatory
Polls					
• Nominal Polling Interval	N/A	Optional*	Mandatory	N/A	Optional[b]
• Tolerated Poll Jitter	N/A	N/A	Optional*	N/A	Optional*

a. N/A means not applicable to this service flow scheduling type. If included in a request for a service flow of this service flow scheduling type, this request MUST be denied.

b. Default is same as Nominal Grant Interval

* Default is CMTS specific

8.2.8 CM Transmit Behavior

In order for these services to function correctly, all that is required of the CM in regards to its transmit behavior for a service flow, is for it to follow the rules specified in Section 7.4.3 and the Request/Transmission Policy specified for the service flow.

8.3 Fragmentation

Fragmentation is an upstream CM "modem capability". The CMTS MUST enable or disable this capability on a per-modem basis with a TLV in the Registration Response. The per-modem basis provides compatibility with DOCSIS1.0 CMs. Once fragmentation is enabled for a DOCSIS 1.1 modem, fragmentation is enabled on a per-Service Flow basis via the Request/Transmission Policy Configuration Settings. When enabled for a Service Flow, fragmentation is initiated by the CMTS when it grants bandwidth to a particular CM with a grant size that is smaller than the corresponding bandwidth request from the CM. This is known as a **Partial Grant**.

8.3.1 CM Fragmentation Support

Fragmentation is essentially encapsulation of a portion of a MAC Frame within a fixed size fragmentation header and a fragment CRC. Concatenated PDUs, as well as single PDUs, are encapsulated in the same manner. Baseline Privacy, if enabled, is performed on each fragment as opposed to the complete original MAC frame.

The CM MUST perform fragmentation according to the flow diagram in Figure 8-8. The phrase "untransmitted portion of packet" in the flow diagram refers to the entire MAC frame when fragmentation has not been initiated and to the remaining untransmitted portion of the original MAC frame when fragmentation has been initiated.

8.3.1.1 Fragmentation Rules:

1. Any time fragmentation is enabled and the grant size is smaller than the request, the CM MUST fill the partial grant it receives with the maximum amount of data (fragment payload) possible accounting for fragmentation overhead and physical layer overhead.
2. The CM MUST send up a piggyback request any time there is no later grant or grant pending for that SID in MAPs that have been received at the CM.
3. If the CM is fragmenting a frame,[6] any piggyback request for the next fragment MUST be made in the BPI EHDR portion of the fragment header. Any piggyback request for a subsequent frame SHOULD be made in the BPI EHDR portion of the last fragment, but MAY be made in one of the extended headers inside the original frame. However, the same request MUST NOT be made in more than one place. Because the CMTS could ignore a request inside the original frame, making the request in the original frame may cause a loss of the request.
4. In calculating bandwidth requests for the remainder of the frame (concatenated frame, if concatenated) that has been fragmented, the CM MUST request enough bandwidth to transmit the entire remainder of the frame plus the 16-byte fragment overhead and all associated physical layer overhead.
5. If the CM does not receive a grant or grant pending within the ACK time of sending a request, the CM MUST backoff and re-request for the untransmitted portion of the frame until the bandwidth is granted or the CM exceeds its retry threshold.
6. If the CM exceeds its retry threshold while requesting bandwidth, the CM discards whatever portion of the frame was not previously transmitted.
7. The CM MUST set the F bit and clear the L bit in the first fragment of a frame.
8. The CM MUST clear the F and L bits in the fragment header for any fragments that occur between the first and last fragments of a frame.
9. The CM MUST set the L bit and clear the F bit in the last fragment of a frame.
10. The CM MUST increment the fragment sequence number sequentially for each fragment of a frame transmitted.

6. Note, 'frame' always refers to either frames with a single Packet PDU or concatenated frames.

11. If a frame is to be encrypted and the frame is fragmented, the frame is encrypted only at the fragment layer with encryption beginning immediately after the fragment header HCS and continuing through the fragment CRC.
12. Frames sent in immediate data (request/data) regions MUST NOT be fragmented.

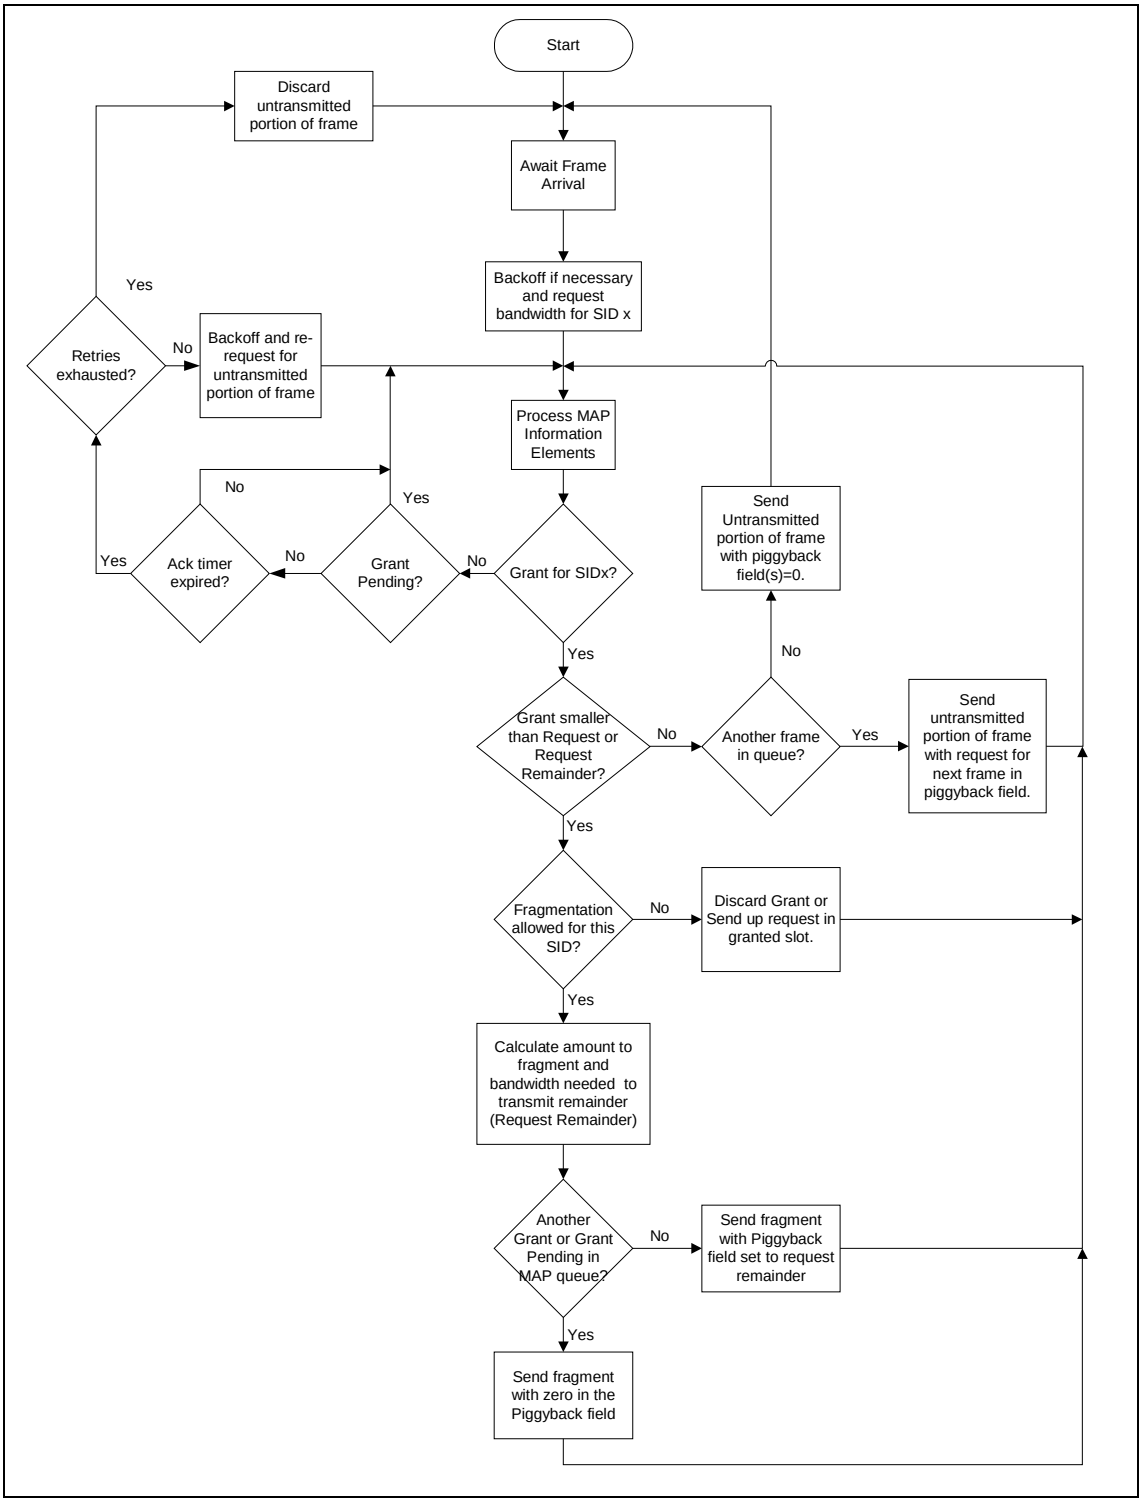


Figure 8-8. CM Fragmentation Flowchart

8.3.2 CMTS Fragmentation Support

At the CMTS, the fragment is processed similarly to an ordinary packet with the exception that the baseline privacy encryption starts just after the fragmentation header as opposed to being offset by 12 bytes.

The CMTS has two modes it can use to perform fragmentation. The Multiple Grant Mode assumes that the CMTS retains the state of the fragmentation. This mode allows the CMTS to have multiple partial grants outstanding for any given SID. The Piggybacking Mode assumes the CMTS does NOT retain any fragmentation state. Only one partial grant is outstanding, so that the CM inserts the remaining amount into the Piggyback field of the fragment header. The type of mode being used is determined by the CMTS. In all cases, the CM operates with a consistent set of rules.

8.3.2.1 Multiple Grant Mode

A CMTS MAY support Multiple Grant Mode for performing fragmentation.

Multiple Grant Mode allows the CMTS to break a request up into two or more grants in a single or over successive maps and it calculates the additional overhead required in the remaining partial grants to satisfy the request. In Multiple Grant Mode, if the CMTS cannot grant the remainder in the current MAP, it MUST send a grant pending (zero length grant) in the current MAP and all subsequent MAPs to the CM until it can grant additional bandwidth. If there is no grant or grant pending in subsequent MAPs, the CM MUST re-request for the remainder. This re-request mechanism is the same as that used when a normal REQ does not receive a grant or grant pending within the ACK time.

If a CM receives a grant pending IE along with a fragment grant, it MUST NOT piggyback a request in the extended header of the fragment transmitted in that grant.

In the case where the CM misses a grant and re-requests the remaining bandwidth, the CMTS MUST recover without dropping the frame.

Due to the imprecision of the mini-slot to byte conversion process the CMTS may not be able to calculate exactly the number of extra mini-slots needed to allow for fragmentation overhead. Also because it is possible for a CM to have missed a map with a partial grant, and thus to be requesting to send an unsent fragment rather than a new PDU, the CMTS can not be certain whether the CM has already accounted for fragmentation overhead in a request. Therefore the CMTS MUST make sure that any fragment payload remainder is at least one mini-slot greater than the number of mini-slots needed to contain the overhead for a fragment (16 bytes) plus the physical layer overhead necessary to transmit a minimum sized fragment. Failure to do this may cause the CMTS to issue a grant that is not needed as the CM has completed transmission of the fragment payload remainder using the previous partial grant. This may cause the CM to get out of sync with the CMTS by inadvertently starting a new fragmentation. Also the CMTS needs to deal with the fact that with certain sets of physical layer parameters, the CM may request one more mini-slot than the maximum size of a short data grant, but not actually need that many mini-slots. This happens in the case where the CM needs to push the request size beyond the short data grant limit. The CMTS needs a policy to ensure that fragmenting such requests in multiple grant mode does not lead to unneeded fragmentary grants.

8.3.2.2 Piggyback Mode

A CMTS MAY support Piggyback Mode for performing fragmentation.

If the CMTS does not put another partial grant or a grant pending in the MAP in which it initiates fragmentation on a SID, the CM MUST automatically piggyback for the remainder. The CM calculates how much of a frame can be sent in the granted bandwidth and forms a fragment to send it. The CM utilizes the piggyback field in the fragment extended header to request the bandwidth necessary to transfer the remainder of the frame. Since the

CMTS did not indicate a multiple grant in the first fragment MAP, the CM MUST keep track of the remainder to send. The request length, including physical-layer and fragmentation overhead, for the remainder of the original frame is inserted into the piggyback request byte in the fragmentation header.

If the fragment HCS is correct, the piggybacked request, if present, is passed on to the bandwidth allocation process while the fragment itself is enqueued for reassembly. Once the complete MAC frame is reassembled and it has been determined that the HCS is correct, the CMTS processes the frame as though it had been received unfragmented except that the CMTS MUST ignore the decryption related portion of any privacy EHDRs. However, the bandwidth requests in privacy EHDRs and request EHDRs of such frame SHOULD be processed, but they MAY be ignored also.

8.3.3 Fragmentation Example

8.3.3.1 Single Packet Fragmentation

Refer to Figure 8-8. Assume that fragmentation has been enabled for a given SID.

1. (Requesting State)- CM wants to transmit a 1018 byte packet. CM calculates how much physical layer overhead (POH) is required and requests the appropriate number of minislots. CM makes a request in a contention region. Go to step 2.
2. (Waiting for Grant)- CM monitors MAPs for a grant or grant pending for this SID. If the CM's ACK time expires before the CM receives a grant or grant pending, the CM retries requesting for the packet until the retry count is exhausted - then the CM gives up on that packet. Go to step 3.
3. (First Fragment)- Prior to giving up in step 2, the CM sees a grant for this SID that is less than the requested number of minislots. The CM calculates how much MAC information can be sent in the granted number of minislots using the specified burst profile. In the example in Figure 8-9, the first grant can hold 900 bytes after subtracting the POH. Since the fragment overhead (FRAG HDR, FHCS, and FCRC) is 16 bytes, 884 bytes of the original packet can be carried in the fragment. The CM creates a fragment composed of the FRAG HDR, FHCS, 884 bytes of the original packet, and an FCRC. The CM marks the fragment as first and prepares to send the fragment. Go to step 4.
4. (First Fragment, multiple grant mode)- CM looks to see if there are any other grants or grant pendings enqueued for this SID. If so, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. Go to step 6. If there are not any grants or grant pendings, go to step 5.
5. (First Fragment, piggyback mode)- If there are no other grants or grant pendings for this SID in this MAP, the CM calculates how many minislots are required to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer for the piggyback request. In the example in Figure 8-9, the CM sends up a request for enough minislots to hold the POH plus 150 bytes (1018-884+16). Go to step 6.
6. (Waiting for Grant)- The CM is now waiting for a grant for the next fragment. If the CM's ACK timer expires while waiting on this grant, the CM should send up a request for enough minislots to send the remainder of the fragmented packet, including the fragmentation overhead, and physical layer overhead. Go to step 7.
7. (Receives next fragment grant)- Prior to giving up in step 6, the CM sees another grant for this SID. The CM checks to see if the grant size is large enough to hold the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead. If so, go to step 10. If not, go to step 8.
8. (Middle Fragment, multiple grant mode)- Since the remainder of the packet (plus overhead) will not fit in the grant, the CM calculates what portion will fit. The CM encapsulates this portion of the packet as a middle fragment. The CM then looks for any other grants or grant pendings enqueued for this SID. If either are

present, the CM sends the fragment with the piggyback field in the FRAG HDR set to zero and awaits the time of the subsequent grant to roll around. Go to step 6. If there are not any grants or grant pendings, go to step 9.

9. (Middle Fragment, piggyback mode) - The CM calculates how many minislots are required to send the remainder of the fragmented packet, including the fragmentation overhead and physical layer overhead, and inserts this amount into the piggyback field of the FRAG HDR. The CM then sends the fragment and starts its ACK timer for the piggyback request. Go to step 6.
10. (Last Fragment) - The CM encapsulates the remainder of the packet as a last fragment. If there is no other packet enqueued or there is a another grant or a grant pending enqueued for this SID, the CM places a zero in the REQ field of the FRAG HDR. If there is another packet enqueued with no grant of grant pending, the CM calculates the number of minislots required to send the next packet and places this number in the REQ field in the FRAG HDR. The CM then transmits the packet. Go to step 11. In the example in Figure 8-9, the grant is large enough to hold the remaining 150 bytes plus POH.
11. (Normal operation)- The CM then returns the normal operation of waiting for grants and requesting for packets. If at any time fragmentation is enabled and a grant arrives that is smaller than the request, the fragmentation process starts again as in step 2.

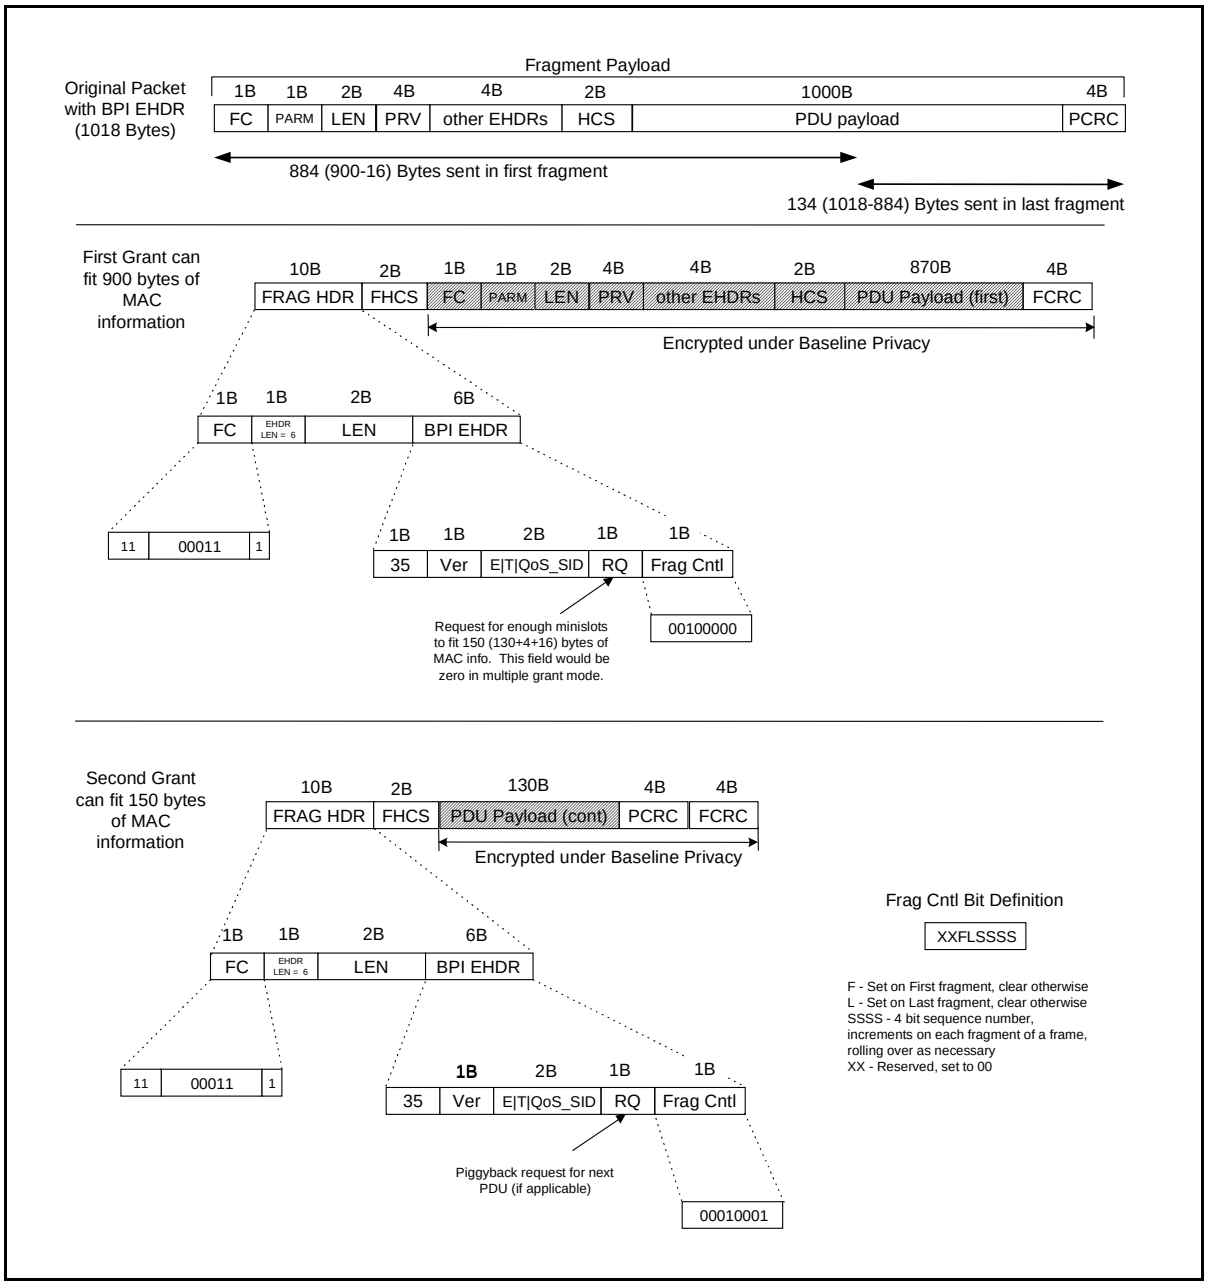


Figure 8-9. Example of Fragmenting a Single Packet

8.3.3.2 Concatenated Packet Fragmentation

After the CM creates the concatenated packet, the CM treats the concatenated packet as a single PDU. Figure 8-10 shows an example of a concatenated packet broken into 3 fragments. Note that the packet is fragmented without regard to the packet boundaries within the concatenated packet.

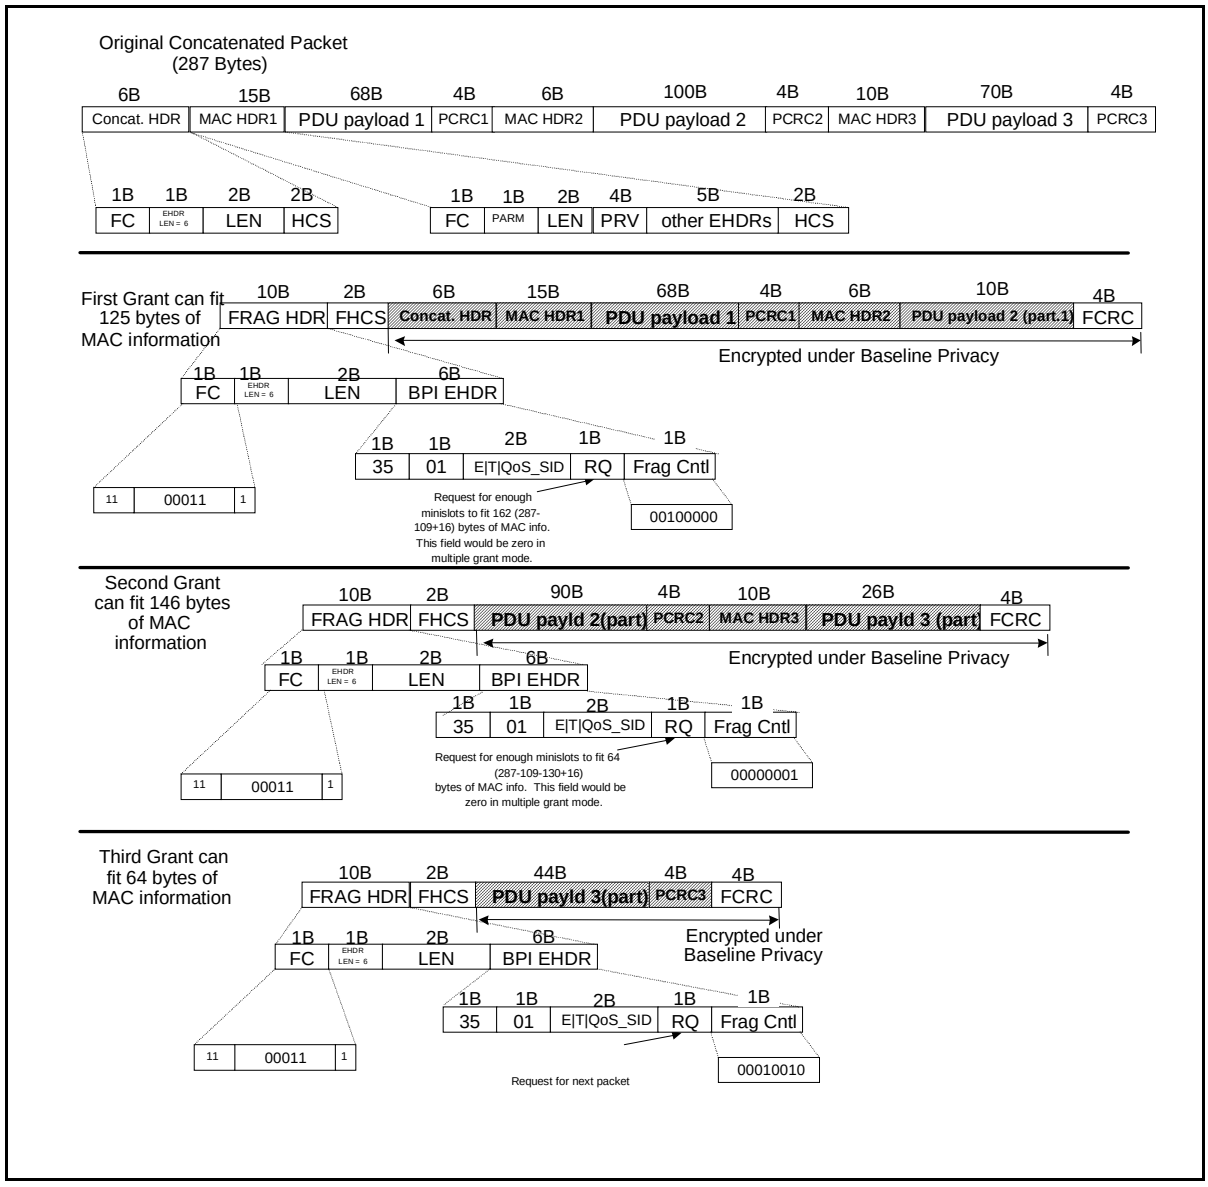


Figure 8-10. Fragmented Concatenated Packet Example

8.4 Payload Header Suppression

The overview section explains the principles of Payload Header Suppression. The subsequent sections explain the signaling for initialization, operation, and termination. Finally, specific upstream and downstream examples are given. The following definitions are used:

Table 8-5. Payload Header Suppression Definitions

PHS	Payload Header Suppression	Suppressing an initial byte string at the sender and restoring the byte string at the receiver.
PHS Rule	Payload Header Suppression Rule	A set of TLV's that apply to a specific PHS Index.
PHSF	Payload Header Suppression Field	A string of bytes representing the header portion of a PDU in which one or more bytes will be suppressed (*i.e.*, a snapshot of the uncompressed PDU header inclusive of suppressed and unsuppressed bytes).
PHSI	Payload Header Suppression Index	An 8-bit value which references the suppressed byte string.
PHSM	Payload Header Suppression Mask	A bit mask which indicates which bytes in the PHSF to suppress, and which bytes to not suppress.
PHSS	Payload Header Suppression Size	The length of the Suppressed Field in bytes. This value is equivalent to the number of bytes in the PHSF and also the number of valid bits in the PHSM
PHSV	Payload Header Suppression Verify	A flag which tells the sending entity to verify all bytes which are to be suppressed.

8.4.1 Overview

In Payload Header Suppression, a repetitive portion of the payload headers following the Extended Header field is suppressed by the sending entity and restored by the receiving entity. In the upstream, the sending entity is the CM and the receiving entity is the CMTS. In the downstream, the sending entity is the CMTS and the receiving entity is the CM. The MAC Extended Header contains a Payload Header Suppression Index (PHSI) which references the Payload Header Suppression Field (PHSF).

Although PHS may be used with any Service Flow Type, it has been designed for use with the Unsolicited Grant Service (UGS) Scheduling Type. UGS works most efficiently with packets of a fixed length. PHS works well with UGS because, unlike other header compression schemes sometimes used with IP data, PHS always suppresses the same number of bytes in each packet. PHS will always produce a fixed length compressed packet header.

The sending entity uses Classifiers to map packets into a Service Flow. The Classifier uniquely maps packets to its associated Payload Header Suppression Rule. The receiving entity uses the Service Identifier (SID) and the PHSI to restore the PHSR.

Once the PHSF and PHSS fields of a rule are known, the rule is considered "fully defined" and none of its fields can be changed. If modified PHS operation is desired for packets classified to the flow, the old rule must be removed from the Service Flow, and a new rule must be installed.

When a classifier is deleted, any associated PHS rule MUST also be deleted.

PHS has a PHSV option to verify or not verify the payload before suppressing it. PHS also has a PHSM option to allow select bytes not to be suppressed. This is used for sending bytes which change such as IP sequence numbers, and still suppressing bytes which do not change.

PHS rules are consistent for all scheduling service types. Requests and grants of bandwidth are specified after suppression has been accounted for. For Unsolicited Grant Services, the grant size is chosen with the Unsolicited Grant Size TLV. The packet with its header suppressed may be equal to or less than the grant size.

The CMTS MUST assign all PHSI values just as it assigns all SID values. Either the sending or the receiving entity MAY specify the PHSF and PHSS. This provision allows for pre-configured headers, or for higher level signaling protocols outside the scope of this specification to establish cache entries. PHS is intended for unicast service, and is not defined for multicast service.

It is the responsibility of the higher-layer service entity to generate a PHS Rule which uniquely identifies the suppressed header within the Service Flow. It is also the responsibility of the higher-layer service entity to guarantee that the byte strings being suppressed are constant from packet to packet for the duration of the Active Service Flow.

8.4.2 Example Applications

- A Classifier on an upstream Service Flow which uniquely defines a Voice-over-IP (VoIP) flow by specifying Protocol Type of UDP, IP SA, IP DA, UDP Source Port, UDP Destination Port, the Service Flow Reference**,** and a PHS Size of 42 bytes. A PHS Rule references this Classifier providing a PHSI value which identifies this VoIP media flow. For the upstream case, 42 bytes of payload header are verified and suppressed, and a 2 byte extended header containing the PHSI is added to every packet in that media flow.
- A Classifier which identifies the packets in a Service Flow, of which 90% match the PHSR. Verification is enabled. This may apply in a packet compression situation where every so often compression resets are done and the header varies. In this example, the scheduling algorithm would allow variable bandwidth, and only 90% of the packets might get their headers suppressed. Since the existence of the PHSI extended header will indicate the choice made, the simple SID/PHSI lookup at the receiving entity will always yield the correct result.
- A Classifier on an upstream Service Flow which identifies all IP packets by specifying Ethertype of IP, the Service Flow ID, a PHSS of 14 bytes, and no verification by the sending entity. In this example, the CMTS has decided to route the packet, and knows that it will not require the first 14 bytes of the Ethernet header, even though some parts such as the Source Address or Destination Address may vary. The CM removes 14 bytes from each upstream frame (Ethernet Header) without verifying their contents and forwards the frame to the Service Flow.

8.4.3 Operation

To clarify operational packet flow, this section describes one potential implementation. CM and CMTS implementations are free to implement Payload Header Suppression in any manner as long as the protocol specified in this section is followed. Figure 8-11 illustrates the following procedure.

A packet is submitted to the CM MAC Service Layer. The CM applies its list of Classifier rules. A match of the rule will result in an Upstream Service Flow, SID, and a PHS Rule. The PHS Rule provides PHSF, PHSI, PHSM, PHSS, and PHSV. If PHSV is set to zero, or is not present, the CM will compare the bytes in the packet header with the bytes in the PHSF that are to be suppressed as indicated by the PHSM. If they match, the CM will suppress all the bytes in the Upstream Suppression Field except the bytes masked by PHSM. The CM will then insert the PHSI into the PHS_Parm field of the Service Flow EH Element, and queue the packet on the Upstream Service Flow.

When the packet is received by the CMTS, the CMTS will determine the associated SID either by internal means or from other Extended Headers elements such as the BPI Extended Header. The CMTS uses the SID and the PHSI to look up PHSF, PHSM, and PHSS. The CMTS reassembles the packet and then proceeds with normal packet processing. The reassembled packet will contain bytes from the PHSF. If verification was enabled, then the PHSF bytes will equal the original header byes. If verification was not enabled, then there is no guarantee that the PHSF bytes will match the original header bytes.

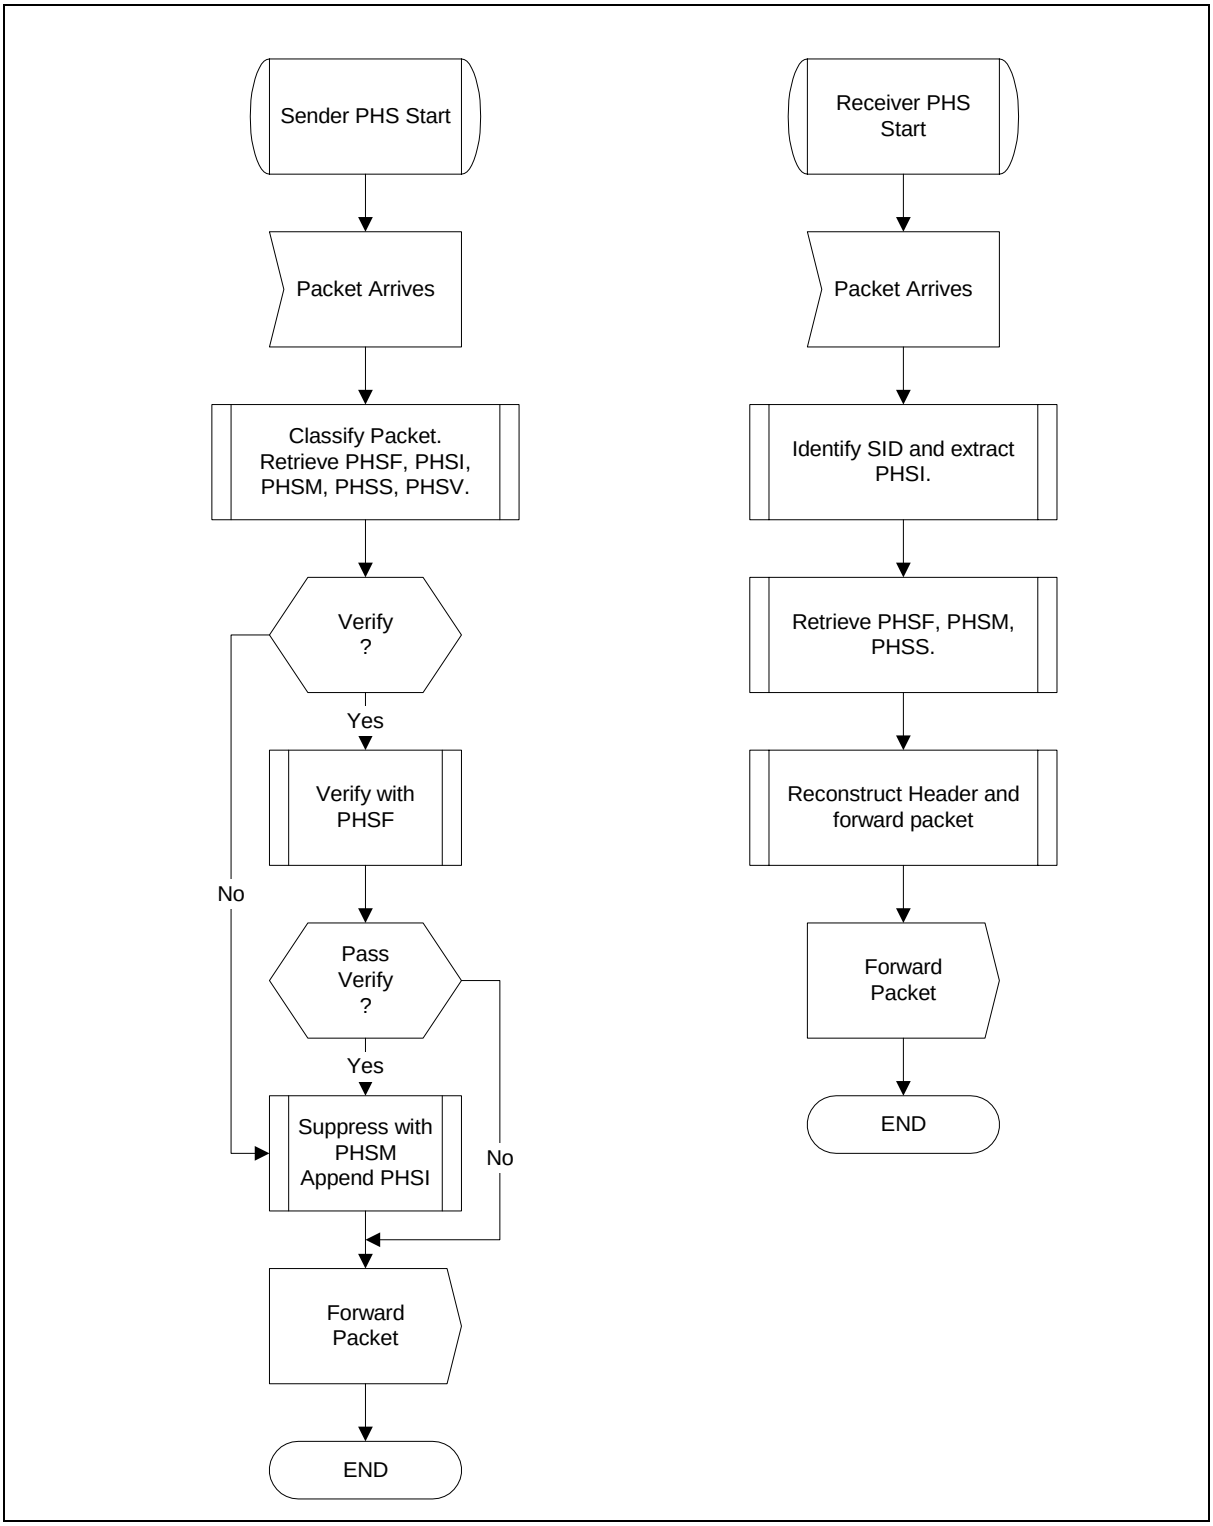


Figure 8-11. Payload Header Suppression Operation

A similar operation occurs in the downstream. The CMTS applies its list of Classifiers. A match of the Classifier will result in a Downstream Service Flow and a PHS Rule. The PHS Rule provides PHSF, PHSI, PHSM, PHSS, and PHSV. If PHSV is set to zero, or is not present, the CMTS will verify the Downstream Suppression Field in the packet with the PHSF. If they match, the CMTS will suppress all the bytes in the Downstream Suppression Field except the bytes masked by PHSM. The CMTS will then insert the PHSI into the PHS_Parm field of the Service Flow EH Element, and queue the packet on the Downstream Service Flow.

The CM will receive the packet based upon the Ethernet Destination Address filtering. The CM then uses the PHSI to lookup PHSF, PHSM, and PHSS. The CM reassembles the packet and then proceeds with normal packet processing.

Figure 8-12 demonstrates packet suppression and restoration when using PHS masking. Masking allows only bytes which do not change to be suppressed. Note that the PHSF and PHSM span the entire Suppression Field, including suppressed and unsuppressed bytes.

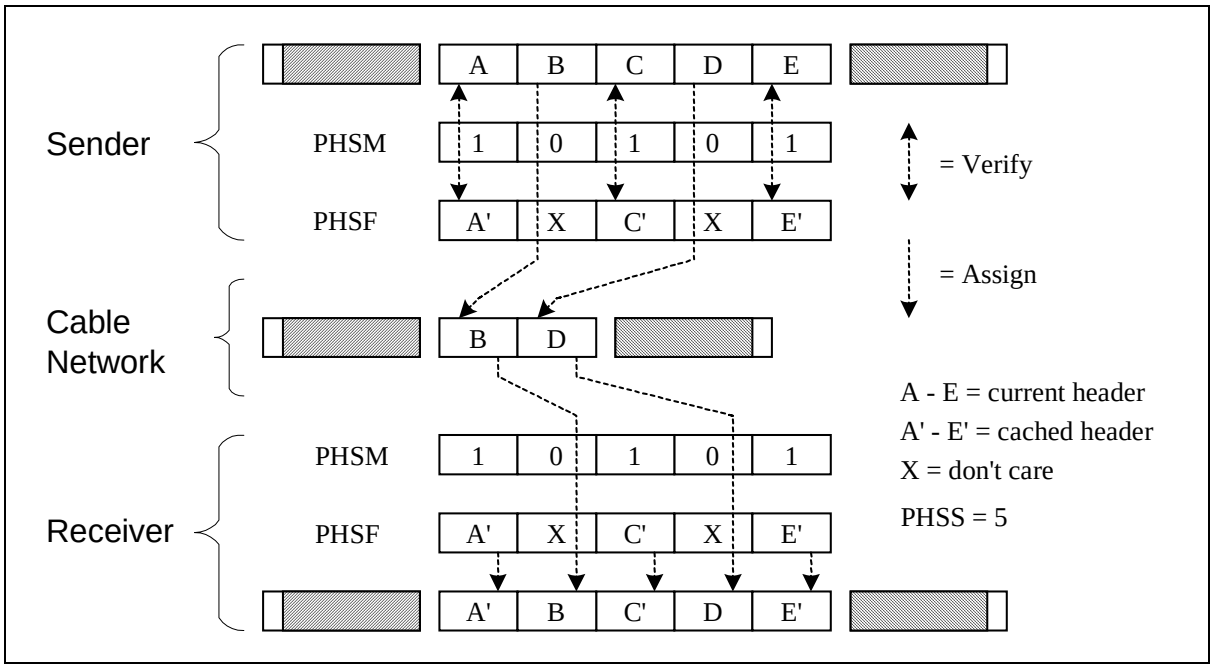


Figure 8-12. Payload Header Suppression with Masking

8.4.4 Signaling

Payload Header Suppression requires the creation of three objects:

- Service Flow
- Classifier
- Payload Header Suppression Rule

These three objects MAY be created in separate message flows, or MAY be created simultaneously.

PHS Rules are created with Registration, DSA, or DSC messages. The CMTS MUST define the PHSI when the PHS Rule is created. PHS Rules are deleted with the DSC or DSD messages. The CM or CMTS MAY define the PHSS and PHSF.

Figure 8-13 shows the two ways to signal the creation of a PHS Rule.

It is possible to partially define a PHS rule (in particular the size of the rule) at the time a Service Flow is created.

As an example, it is likely that when a Service Flow is first provisioned the size of the header field to be suppressed will be known. The values of some items within the field (*e.g.*, IP addresses, UDP port numbers, etc.) may not be known and would be provided in a subsequent DSC as part of the activation of the Service Flow (using the "Set PHS Rule" DSC Action).

A PHS rule is partially defined when the PHSF and PHSS field values are not both known. Once both PHSF and PHSS are known, the rule is considered fully defined, and MUST NOT be modified via DSC signaling. PHSV and PHSM fields have default values, thus are not required to fully define a PHS rule. If PHSV and PHSM are not known when the rule becomes fully defined, their default values are used, and MUST NOT be modified via DSC signaling.

Each step of the PHS rule definition, whether it is a registration request, DSA or a DSC, MUST contain Service Flow ID (or reference), Classifier ID (or reference) to uniquely identify the PHS rule being defined. A PHS Index and Service ID pair is used to uniquely identify the PHS rule during upstream packet transfer. A PHS Index is enough to uniquely identify the PHS rule used in downstream packet transfer.

8.4.5 Payload Header Suppression Examples

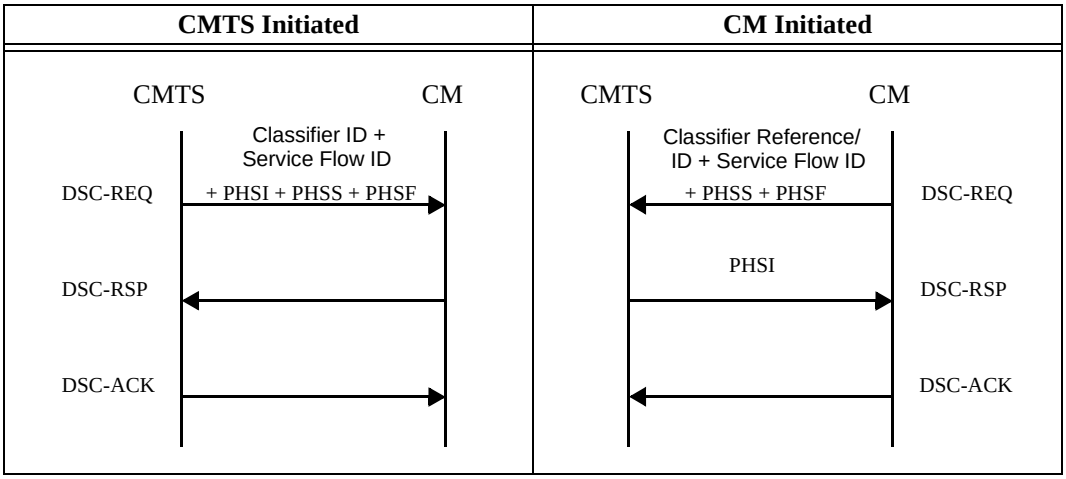


Figure 8-13. Payload Header Suppression Signaling Example

8.4.5.1 Upstream Example

A Service Class with the Service Class Name of "G711-US-UGS-HS-42" is established which is intended for ITU-T Recommendation G.711 VoIP traffic in the upstream with Unsolicited Grant Service. When Classifiers are added to the flow, a PHSS value of 42 is included which explicitly states that the first 42 bytes following the MAC Extended Header on all packets in that flow must be verified, suppressed, and restored. In this example, the Service Class is configured such that any packet which does not verify correctly will not have its header suppressed and will be discarded since it will exceed the Unsolicited Grant Size. (Refer to Appendix C.2.2.6.3)

Figure 8-14 shows the encapsulation used in the upstream with and without Payload Header Suppression. An RTP Voice over IP Payload without IPsec is used as a specific example to demonstrate efficiency.

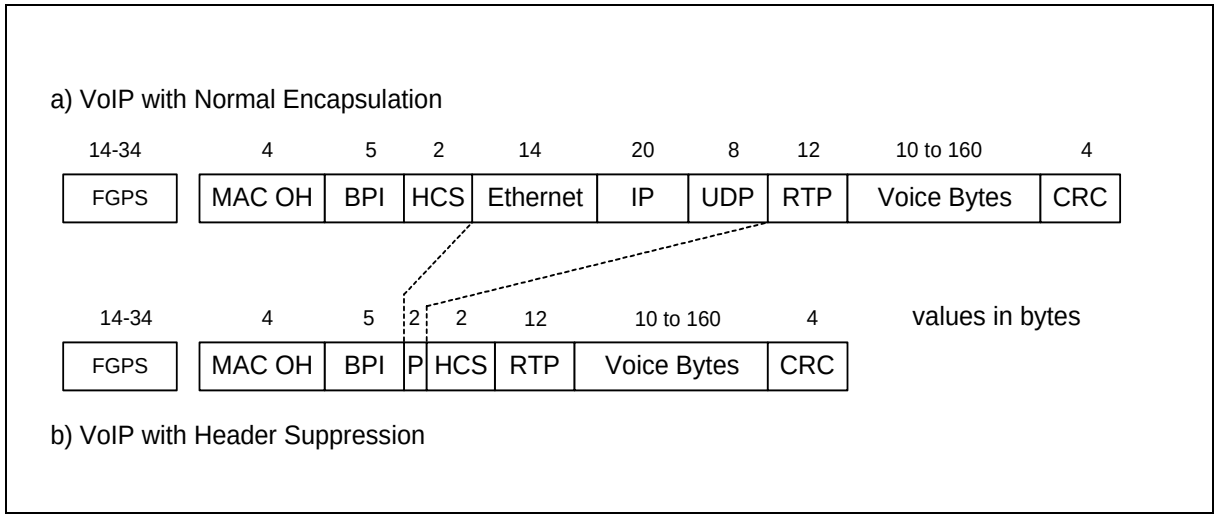


Figure 8-14. Upstream Payload Header Suppression Example

Figure 8-14a shows a normal RTP packet carried on an upstream channel. The beginning of the frame represents the physical layer overhead (FGPS) of FEC, guard time, preamble, and stuffing bytes. Stuffing bytes occur in the last code word and when mapping blocks to minislots. Next is the MAC layer overhead including the 6 byte MAC header with a 5 byte BPI Extended Header, the 14 byte Ethernet Header, and the 4 byte Ethernet CRC trailer. The VoIP payload uses a 20 byte IP header, an 8 byte UDP header, and a 12 byte RTP header. The voice payload is variable and depends upon the sample time and the compression algorithm used.

Figure 8-14b shows the same payload with Payload Header Suppression enabled. In the upstream, Payload Header Suppression begins with the first byte after the MAC Header Checksum. The 14 byte Ethernet header, the 20 byte IP header, and the 8 byte UDP header have been suppressed, and a 2 byte PHS Extended Header element has been added, for a net reduction of 40 bytes. In this example of an established VoIP connection, these fields remain constant from packet to packet, and are otherwise redundant.

8.4.5.2 Downstream Example

A Service Class with the Service Class Name of "G711-DS-HS-30" is established which is intended for G.711 VoIP traffic in the downstream. When Classifiers are added to the Service Flow, a PHSS value of 30 is included which explicitly indicates that 30 bytes of the payload header on all packets must be processed for suppression and restoration according to the PHSM. Any packet which does not verify correctly will not have its header suppressed but will be transmitted subject to the traffic shaping rules in place for that Service Flow.

Figure 8-15 shows the encapsulation used in the downstream with and without Payload Header Suppression. An RTP Voice over IP Payload without IPsec is used as a specific example to demonstrate efficiency.

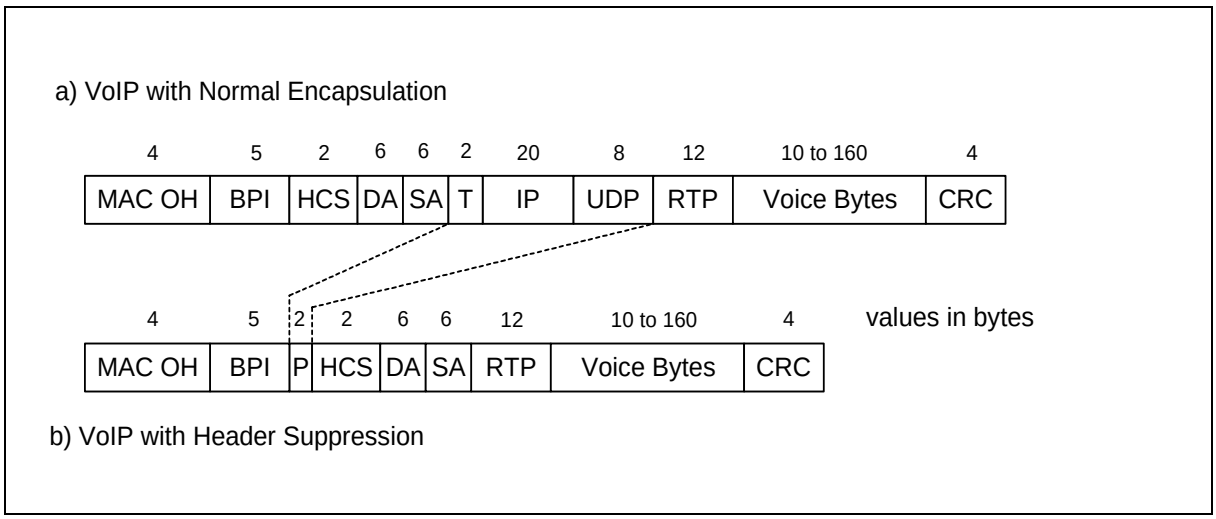


Figure 8-15. Downstream Payload Header Suppression Example

Figure 8-15a shows a normal RTP packet carried on a downstream channel. The Layer 2 overhead includes the 6 byte MAC header with a 5 byte BPI Extended Header, the 14 byte Ethernet Header (6 byte Destination Address, 6 byte Source Address, and 2 byte EtherType field), and the 4 byte Ethernet CRC trailer. The Layer 3 VoIP payload uses a 20 byte IP header, an 8 byte UDP header, and a 12 byte RTP header. The voice payload is variable and depends upon the sample time and the compression algorithm used.

Figure 8-15b shows the same payload with Payload Header Suppression enabled. In the downstream, Payload Header Suppression begins with the thirteenth byte after the MAC Header Checksum. This retains the Ethernet Destination Address and Source Address which is required so that the CM may filter and receive the packet. The remaining 2 bytes of the Ethernet Header, the 20 byte IP header, and the 8 byte UDP header have been suppressed, and a 2 byte PHS Extended Header element has been added, for a net reduction of 28 bytes. In this example of an established VoIP connection, these fields remain constant from packet to packet, and are thus redundant.

9 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.

9.1 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.

9.2 Cable Modem Initialization

The procedure for initialization of a cable modem MUST be as shown in Figure 9-1. 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 Appendix 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
- Device Class Identification (optional)
- Establish IP connectivity
- Establish time of day
- Transfer operational parameters
- 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] (*e.g.,* X.509 certificate) used to authenticate the CM to the security server and authenticate the responses from the security and provisioning servers.

The SDL (Specification and Description Language) notation used in the following figures is shown in Figure 9-2 (refer to ITU-T Recommendation Z.100 [ITU-T Z.100]).

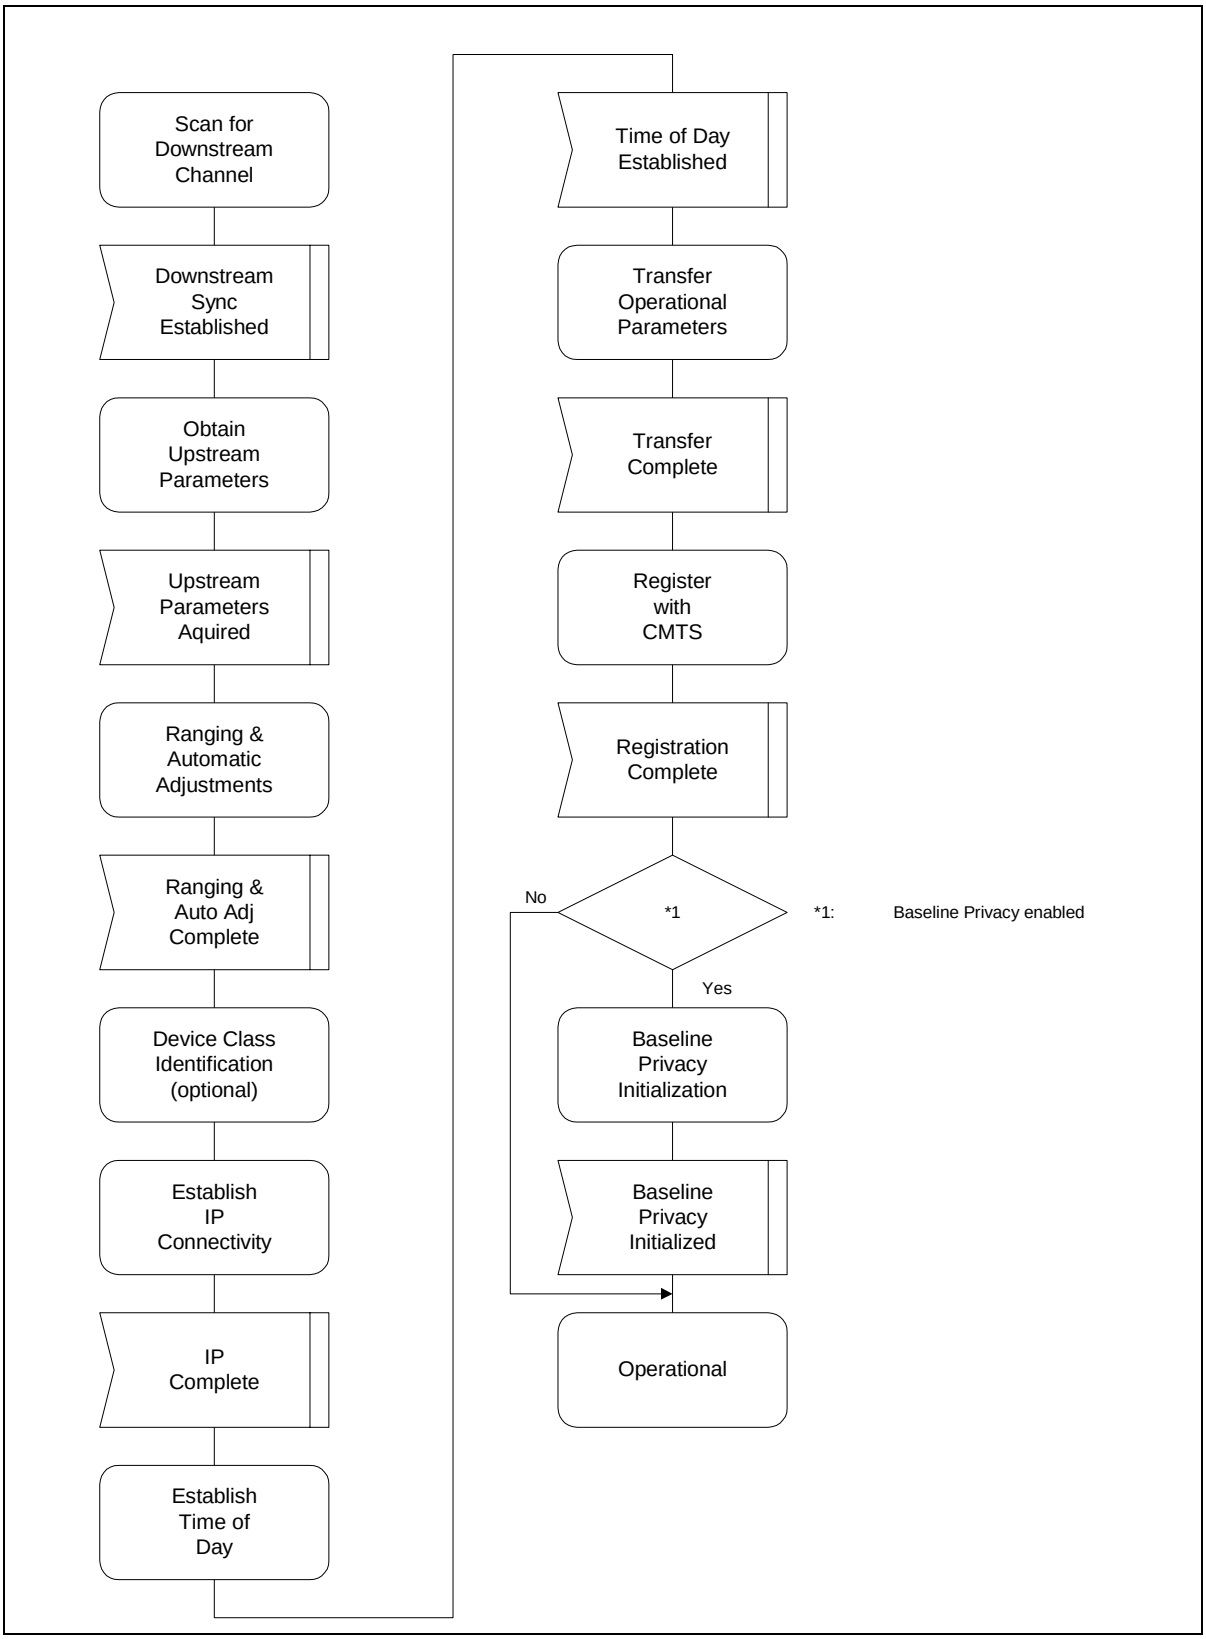


Figure 9-1. CM Initialization Overview

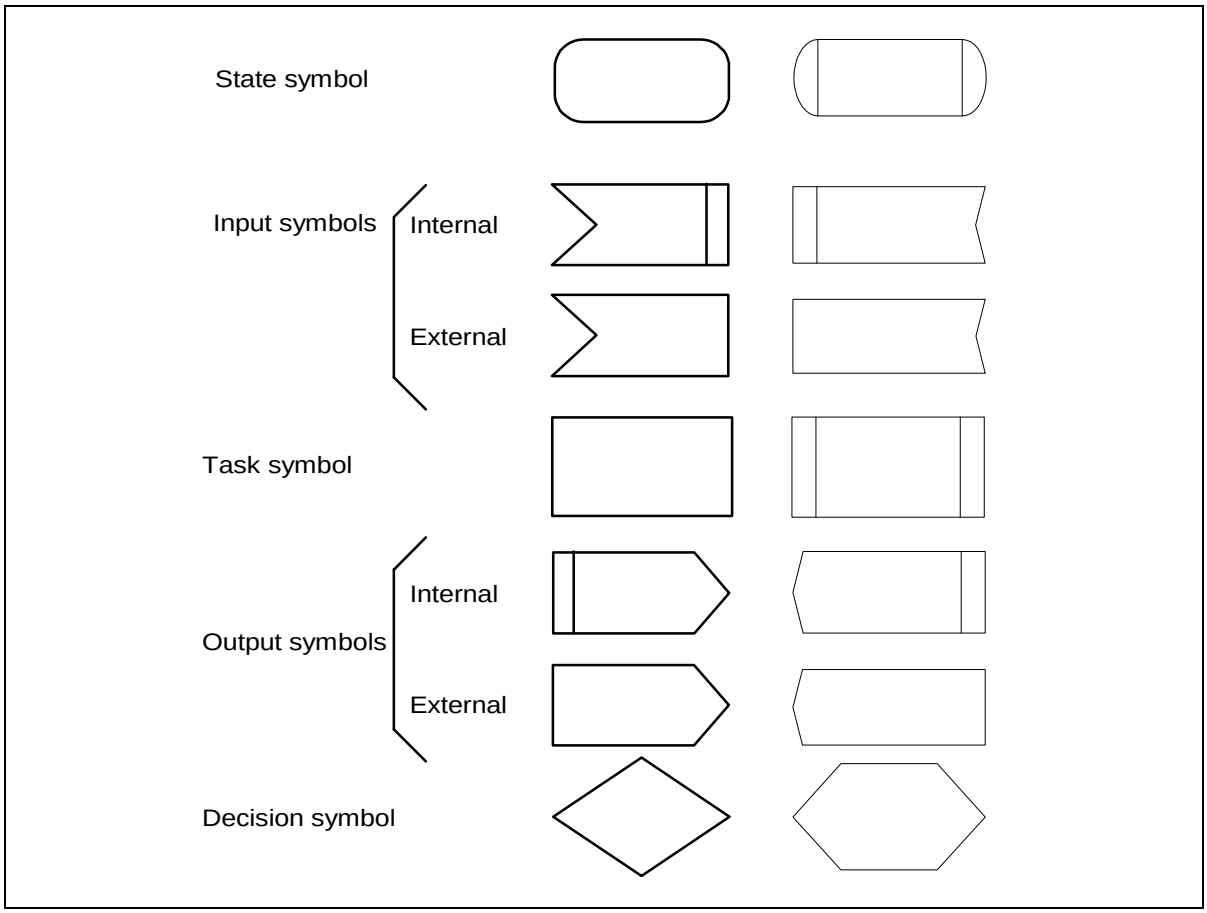


Figure 9-2. SDL Notation

9.2.1 Scanning and Synchronization to Downstream

On initialization or a "Reinitialize MAC" operation, 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.

In order to support redundant CMTS architectures, when a CM in the Operational state detects that the downstream signal is invalid (*i.e.*, does not meet the four criteria above), the CM MUST NOT immediately perform a Reinitialize MAC operation. It must instead attempt to re-establish synchronization on the current downstream channel (see Section 9.5). Such re-establishment attempts MUST continue until the operation of Periodic Ranging as specified in Figure 9-17 of Section 9.3.1calls for a "Re-initialize MAC" operation after the

expiration of Timeout T4 or 16 expirations of Timeout T3. Figure 9-17 shows the procedure that MUST be followed by a cable modem during standard operation.

9.2.2 Obtain Upstream Parameters

Refer to Figure 9-3. 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 usable channel IDs. If no channel can be found after a suitable timeout period, then 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, then the CM MUST try the next channel ID until it finds a usable channel. If the channel is suitable, the CM MUST extract the parameters for this upstream from the UCD. It then MUST wait for the next SYNC message[1] 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 9-6. If initial ranging is not successful, then the next channel ID is selected, and the procedure restarted from UCD extraction. When there are no more channel IDs to try, then the CM MUST continue scanning to find another downstream channel.

1. 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.

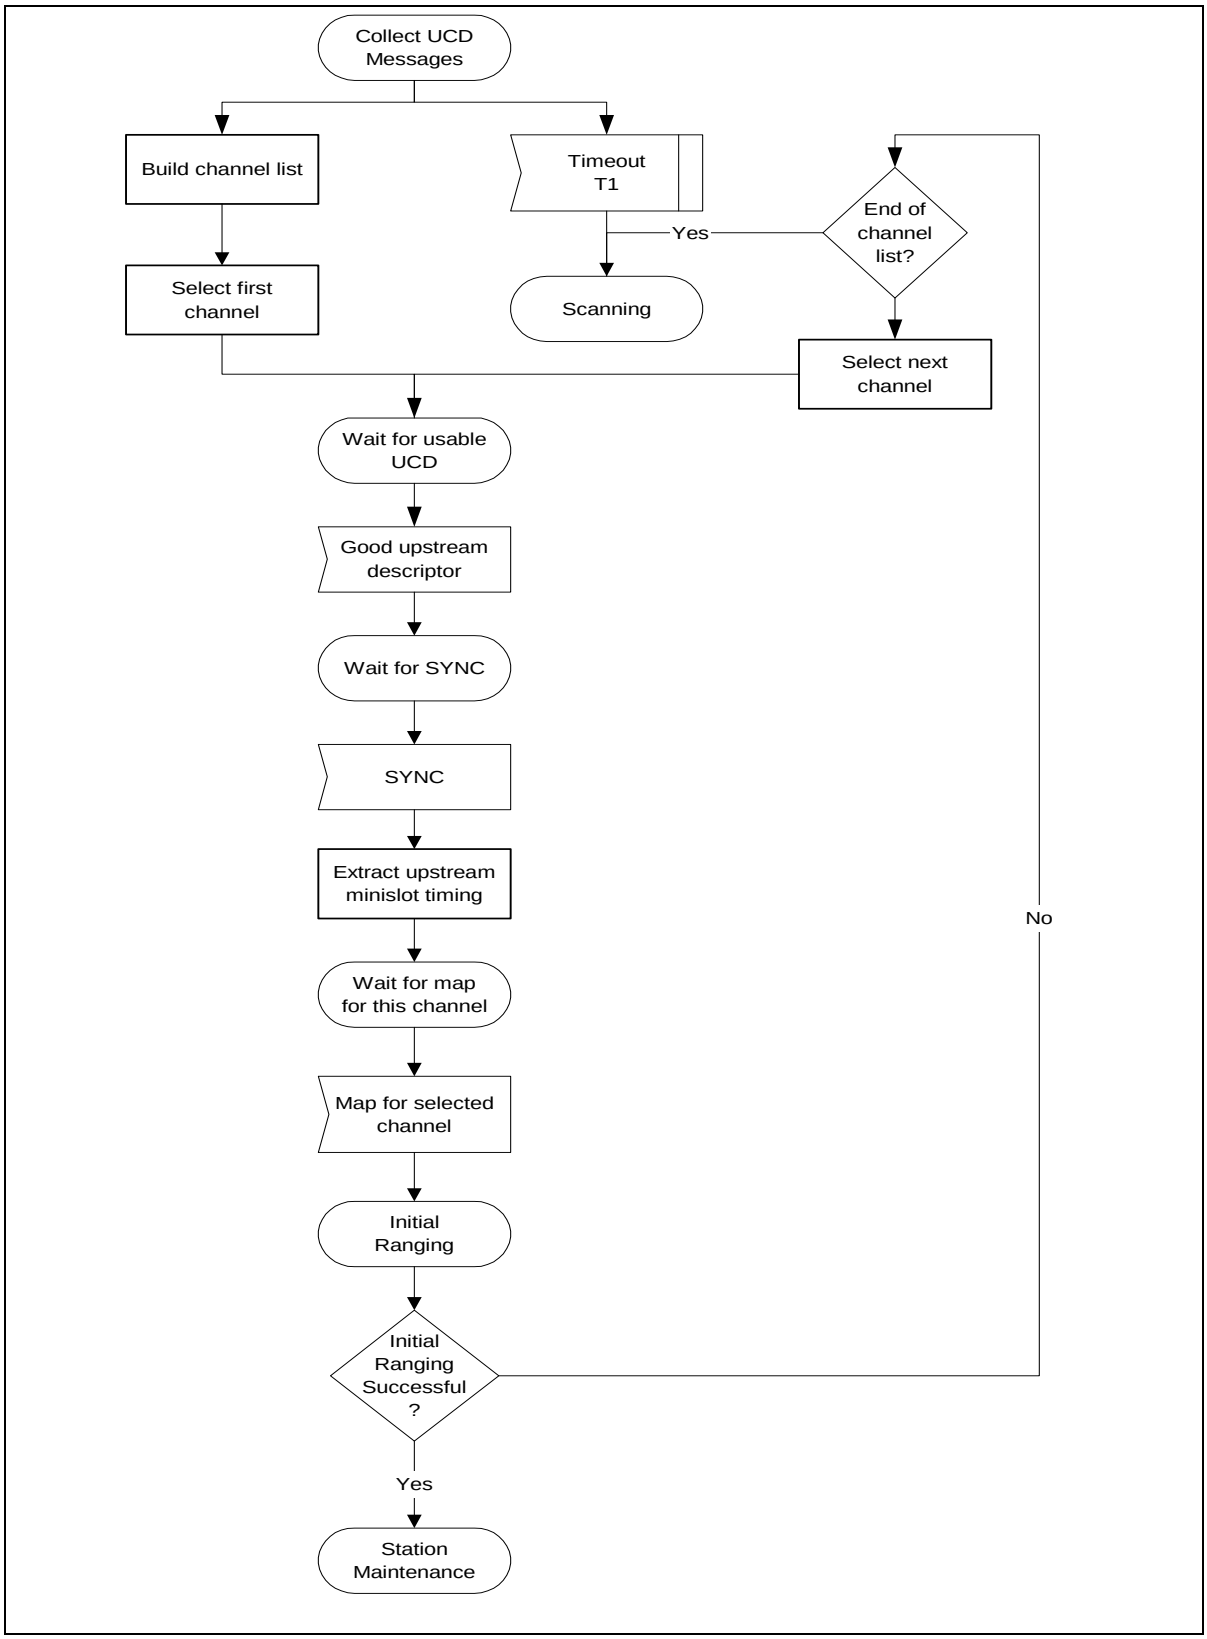


Figure 9-3. Obtaining Upstream Parameters

9.2.3 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 Appendix B. These messages are addressed to all CMs. Refer to Figure 9-4.

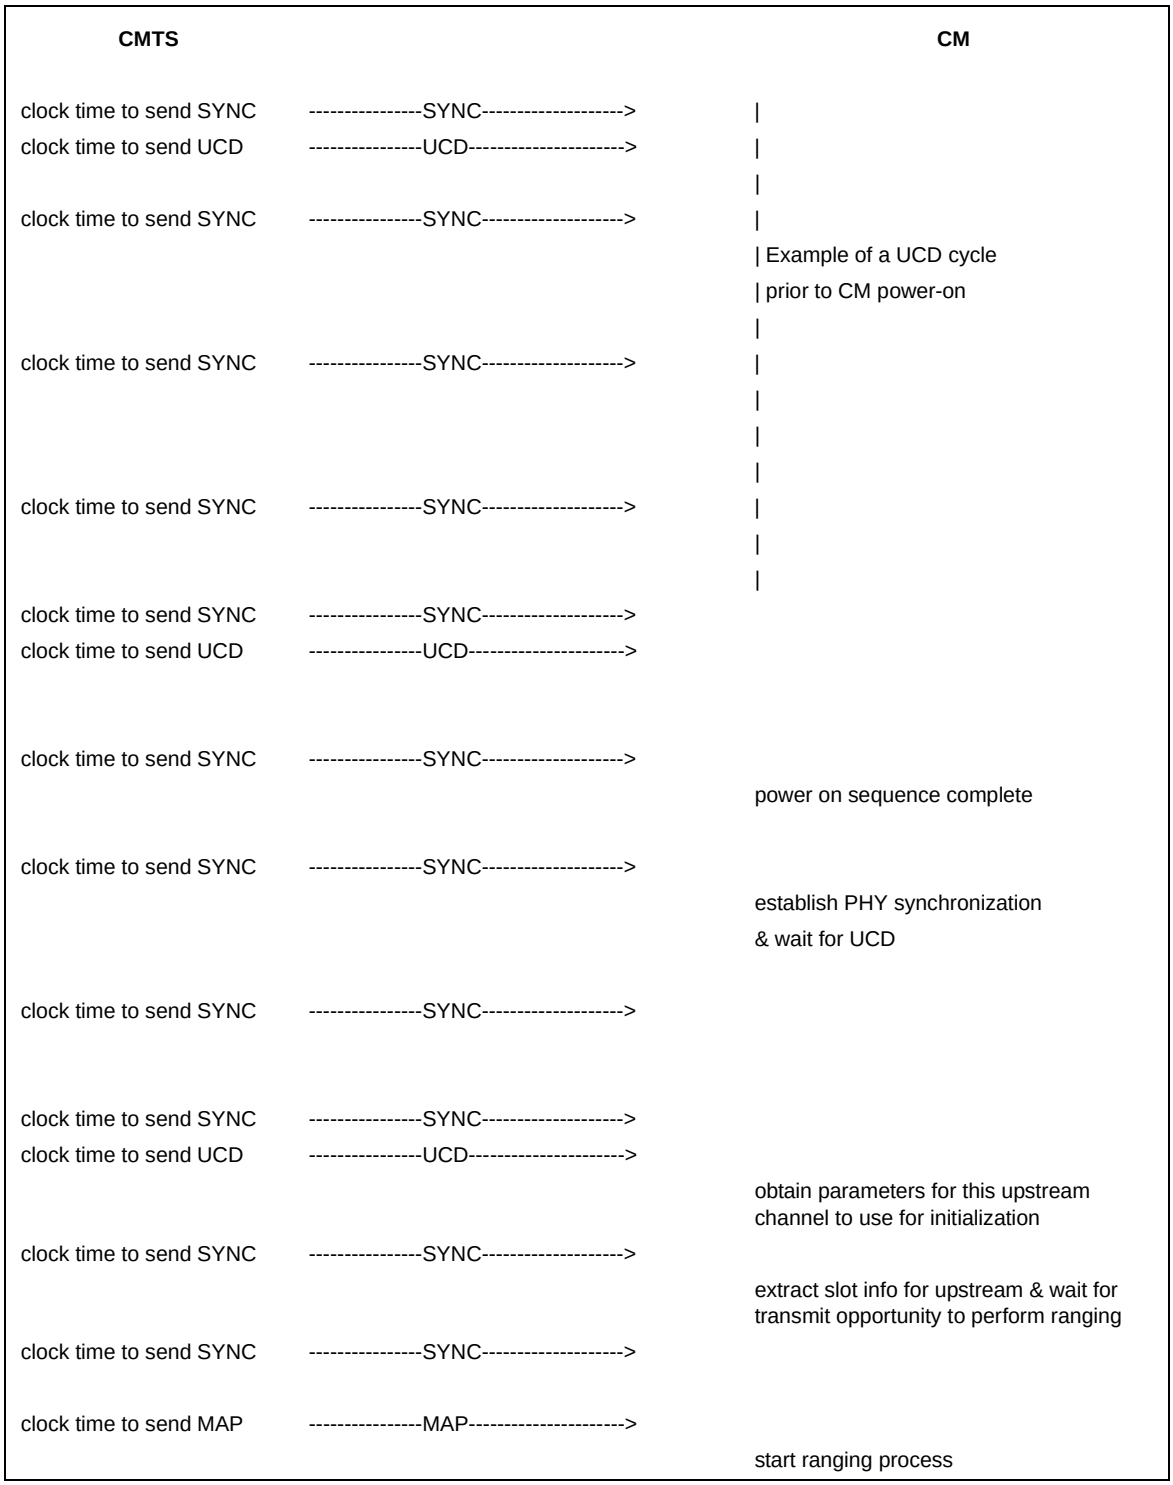


Figure 9-4. Message Flows During Scanning and Upstream Parameter Acquisition

9.2.4 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 CMTSs. Refer to Figure 9-5 through Figure 9-8.

Note: MAPs are transmitted as described in Section 6.

CMTS		CM
[time to send the Initial Maintenance opportunity]		
send map containing Initial Maintenance information element with a broadcast/multicast Service ID	----------MAP---------->	
	<---------RNG-REQ-------	transmit ranging packet in contention mode with Service ID parameter = 0
[receive recognizable ranging packet]		
allocate temporary Service ID		
send ranging response	----------RNG-RSP------->	
add temporary Service ID to poll list		store temporary Service ID & adjust other parameters
[time to send the next map]		
send map with Station Maintenance information element to modem using temporary SID	----------MAP---------->	recognize own temporary Service ID in map
	<---------RNG-REQ-------	reply to Station Maintenance opportunity poll
send ranging response	----------RNG-RSP------->	
		adjust local parameters
[time to send an Initial Maintenance opportunity] send map containing Initial Maintenance information element with a broadcast/multicast Service ID		
send periodic transmit opportunity to broadcast address	----------MAP---------->	

Figure 9-5. 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 Appendix B.

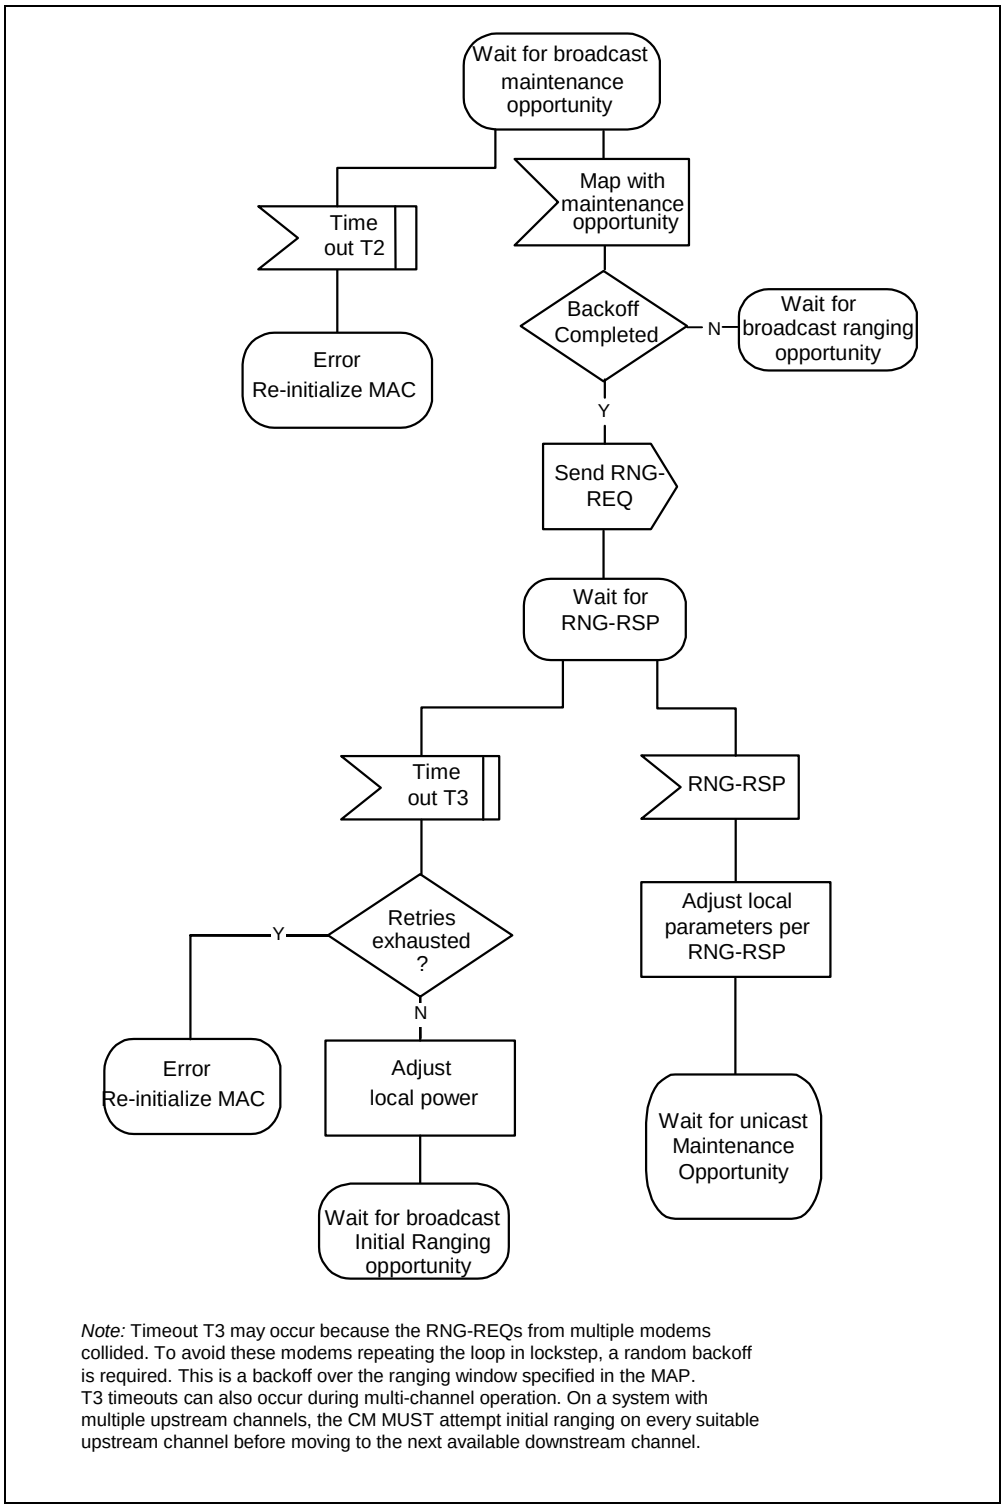
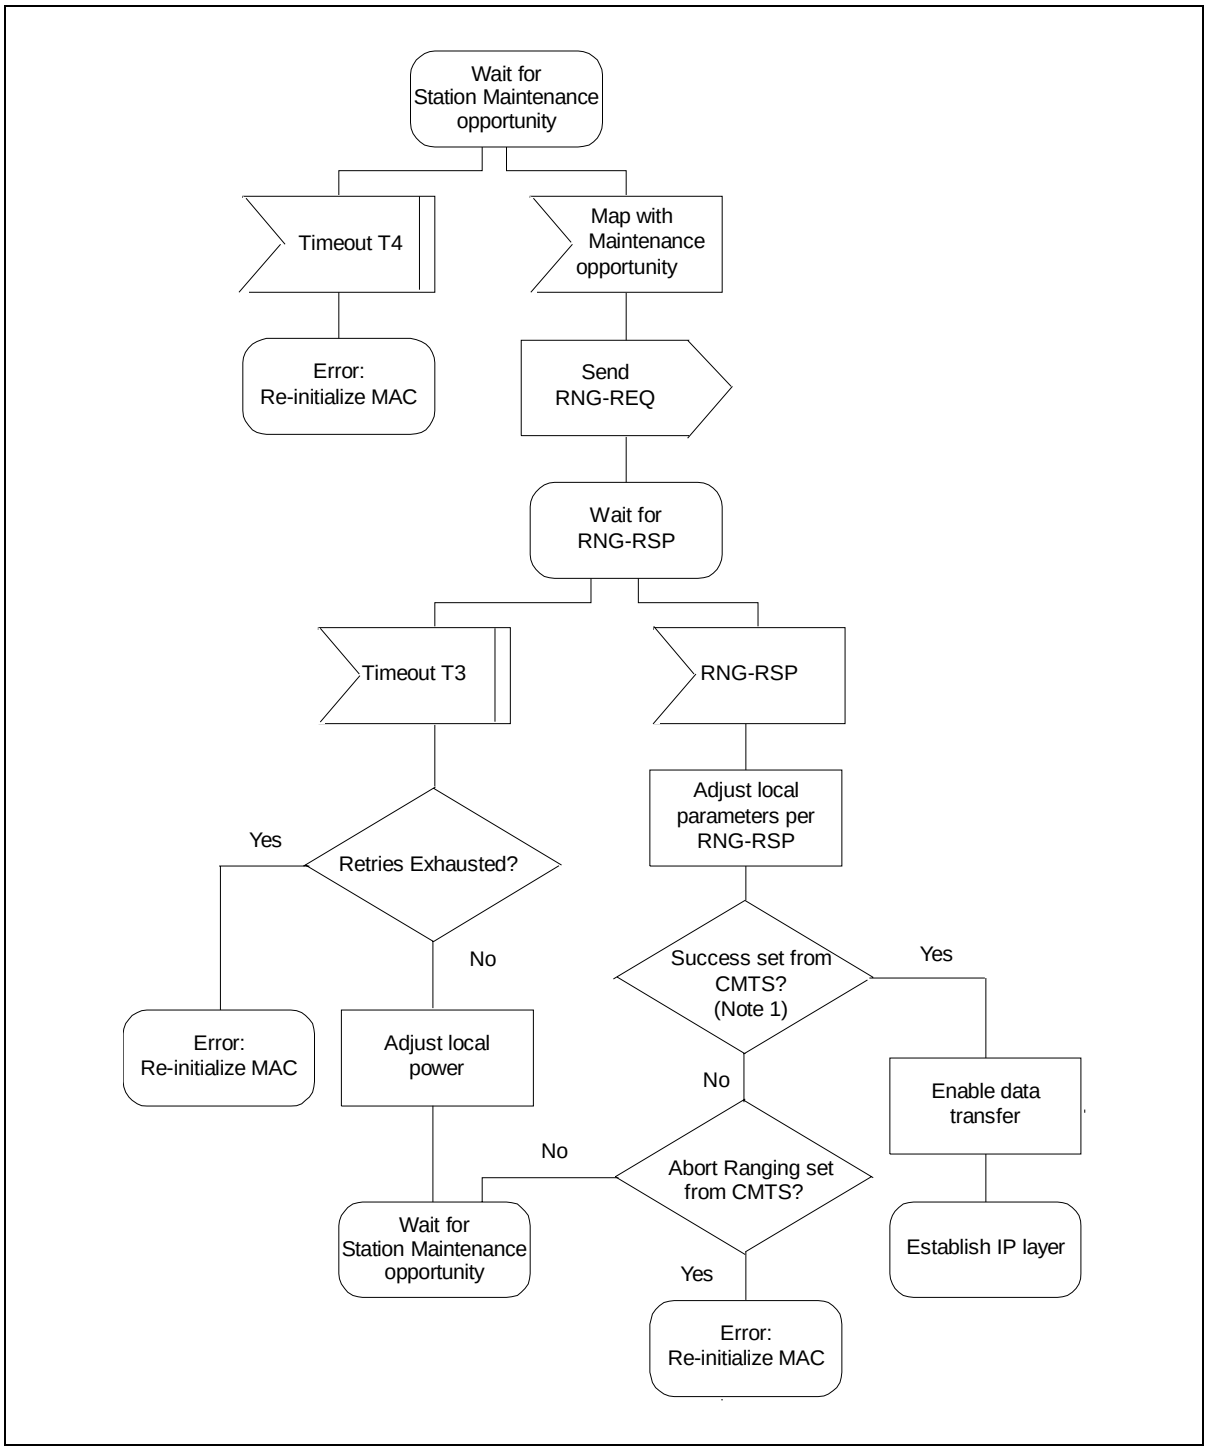


Figure 9-6. Initial Ranging - CM



1. Ranging Request is within the tolerance of the CMTS.

Figure 9-7. Initial Ranging - CM (continued)

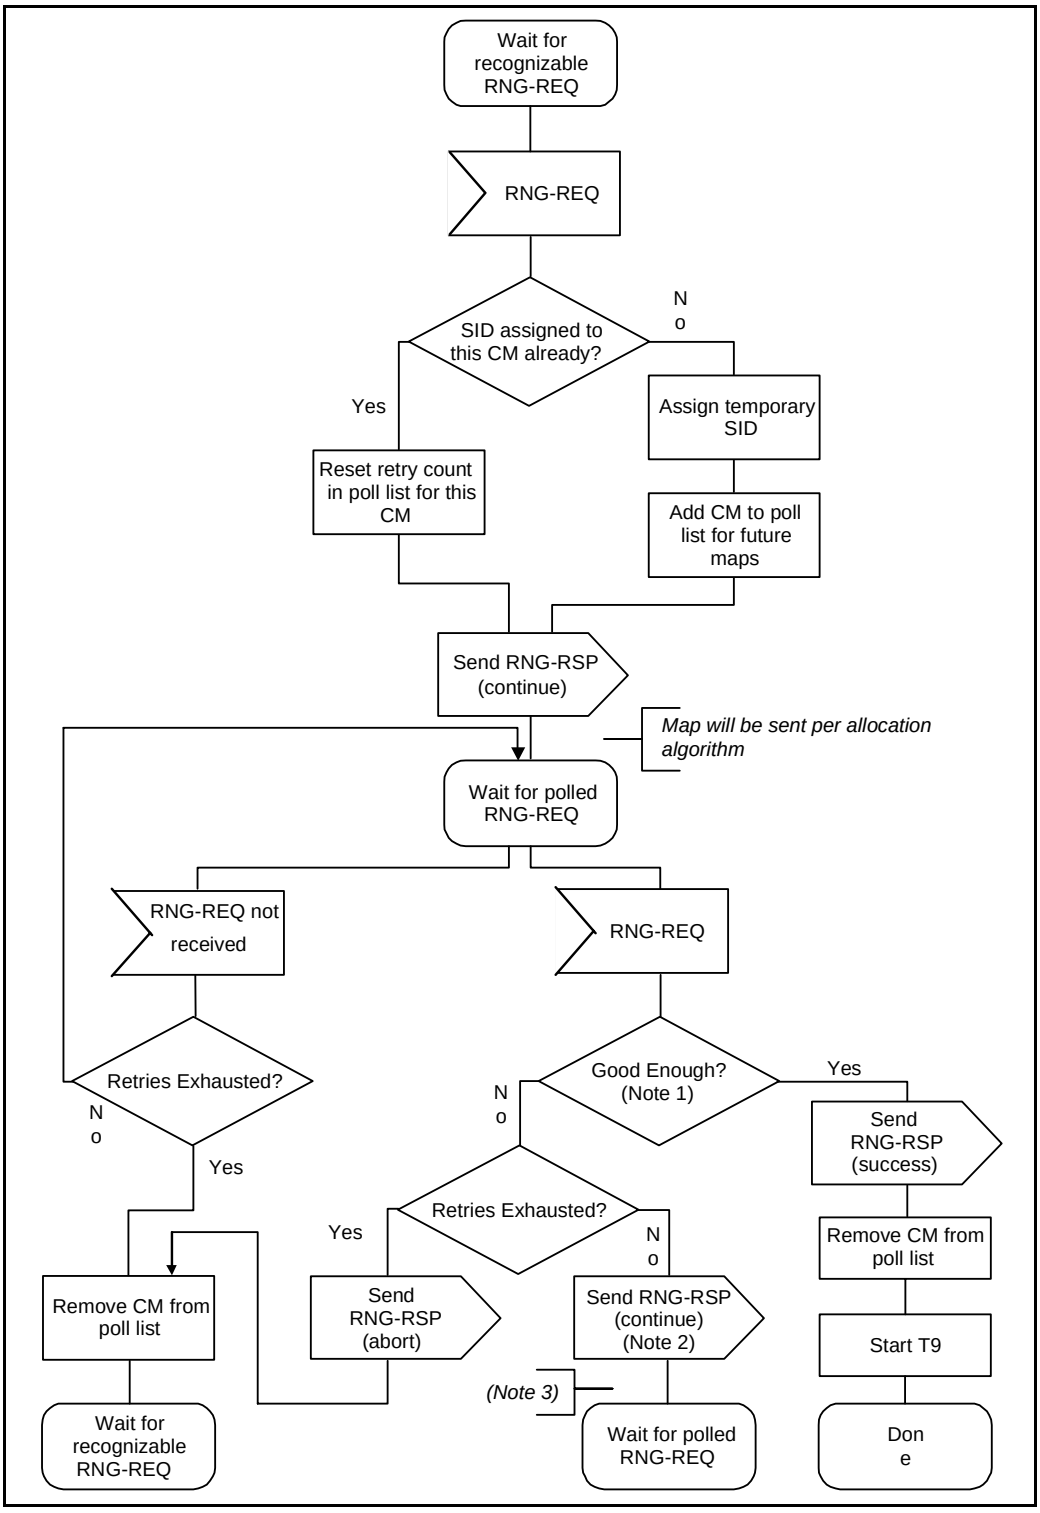


Figure 9-8. Initial Ranging - CMTS

Notes to Figure 9-8:

1. Means pending-till-complete was zero and ranging is within the tolerable limits of the CMTS.
2. If pending-till-complete is nonzero, the CMTS MUST NOT send new Pre-Equalization Coefficients in the corresponding RNG-RSP. However, the CMTS MAY adjust other ranging parameters in the RNG-RSP message.
3. If the RNG-REQ pending-till-complete was nonzero, the CMTS MAY schedule additional station maintenance opportunities during the pending-till-complete period in order to adjust ranging parameters other than the equalizer.

9.2.4.1 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 6.3.6):

- 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 increased to the maximum value (without a response from the CMTS) 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 6.3.6 and in Appendix H.

9.2.5 Device Class Identification

After Ranging is complete and before establishing IP connectivity, the CM MAY identify itself to the CMTS for use in provisioning. Refer to Figure 9-9.

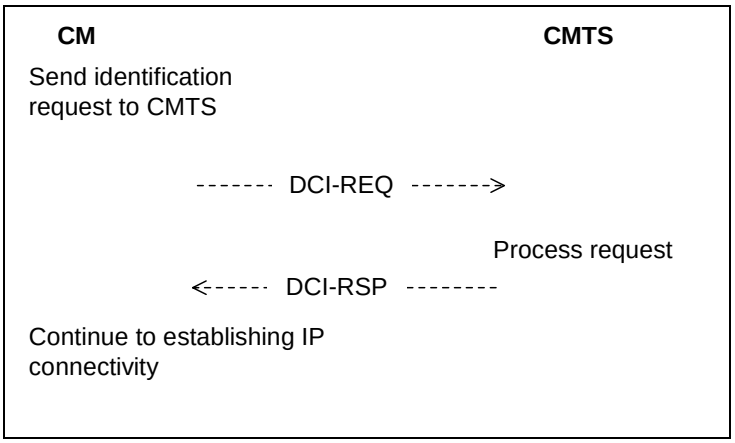


Figure 9-9. Device Class Identification

If implemented, the CM MUST use an adaptive timeout for device class identification based on binary exponential backoff, similar to that used for TFTP. Refer to Section 9.2.8 for details.

9.2.6 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 (refer to Appendix D). The DHCP response MUST contain the name of a file which contains further configuration parameters. Refer to Figure 9-10.

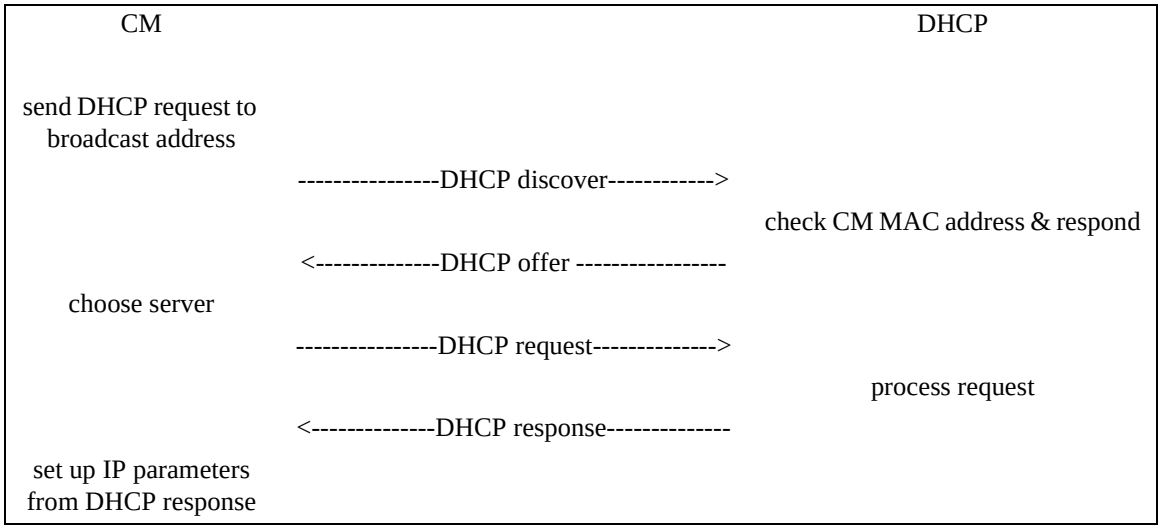


Figure 9-10. Establishing IP Connectivity

9.2.7 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 retrieved 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 9-11. 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.

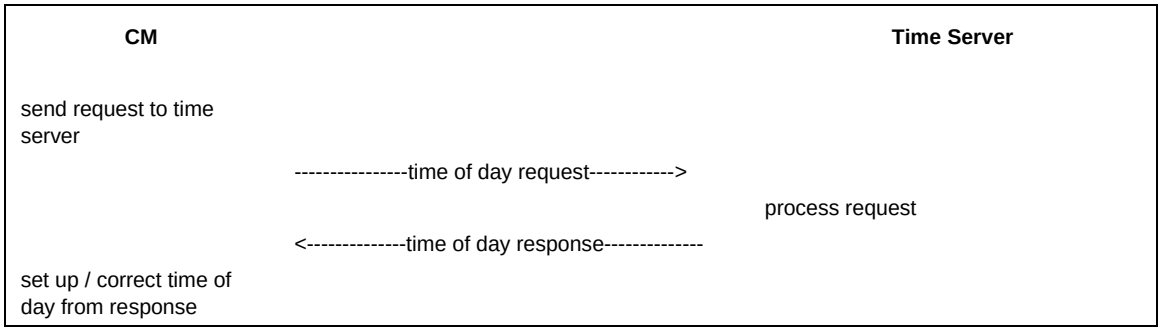


Figure 9-11. Establishing Time of Day

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.

9.2.8 Transfer Operational Parameters

After DHCP is successful, the modem MUST download the parameter file using TFTP, as shown in Figure 9-12. 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].

The parameter fields required in the DHCP response and the format and content of the configuration file MUST be as defined in Appendix D. 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 6.3.6.3. The modem MAY reject the configuration file in the case of size limit errors (refer to Appendix D.2.1).

9.2.9 Registration

A CM MUST be authorized to forward traffic into the network once it is initialized and configured. The CM is authorized to forward traffic into the network via registration. To register with a CMTS, the CM MUST forward its configured class of service and any other operational parameters in the configuration file (refer to Section 6.3.7) to the CMTS as part of a Registration Request. The CM MUST perform the following operations before registration (refer to Figure 9-12):

- Check the mandatory items in the configuration file (refer to Appendix D.2.2). The CM MUST reject the configuration file if it lacks any mandatory items.
- Calculate a MIC per Appendix D.2.3.1 and compare it to the CM MIC included in the configuration file. If the MIC is invalid, the CM MUST reject the configuration file.
- If the configuration file contains TLV-11 encoding, the CM MUST follow the configuration process defined in [DOCSIS5] Section 3.4. The CM MUST reject the configuration file in the case of TLV-11 processing failure.

The configuration parameters downloaded to the CM MUST include a network access control object (see Appendix C.1.1.3). If this is set to "no forwarding," the CM MUST NOT forward data from attached CPE to the network, yet the CM MUST respond to network management requests. This allows the CM to be configured in a mode in which it is manageable but will not forward data.

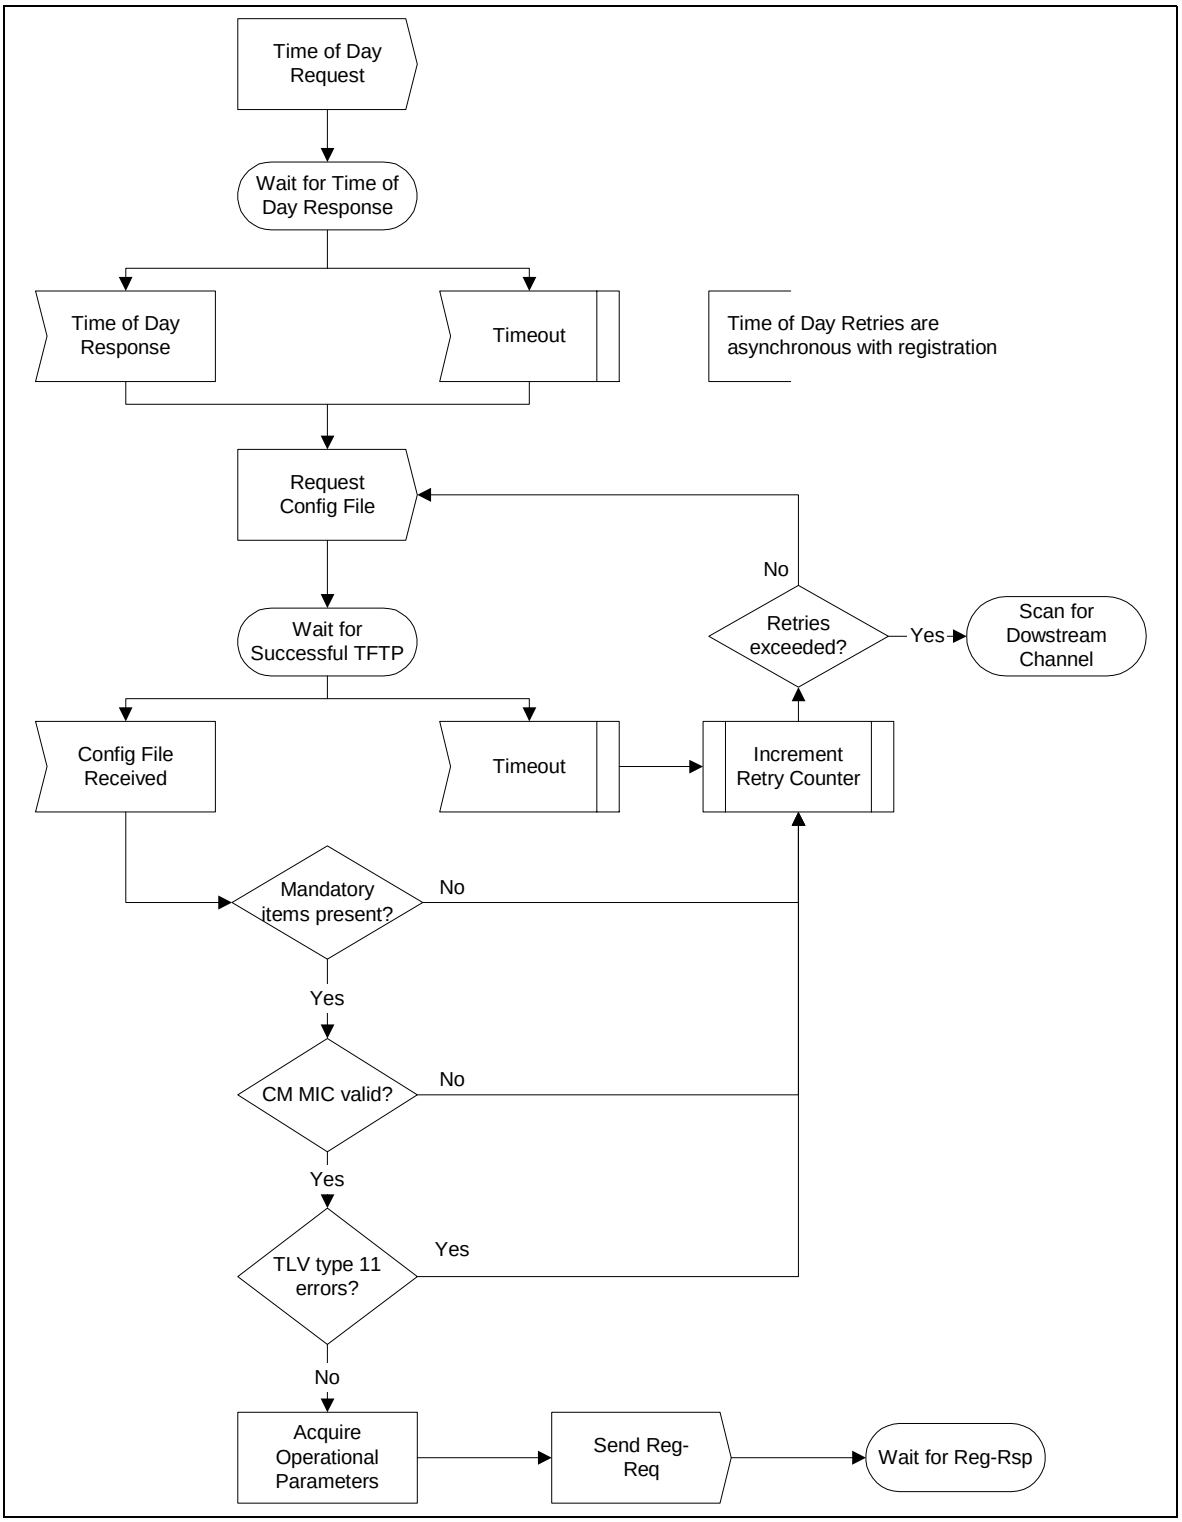


Figure 9-12. Registration — CM

Once the CM has sent a Registration Request to the CMTS it MUST wait for a Registration Response to authorize it to forward traffic to the network. Figure 9-13 shows the waiting procedure that MUST be followed by the CM.

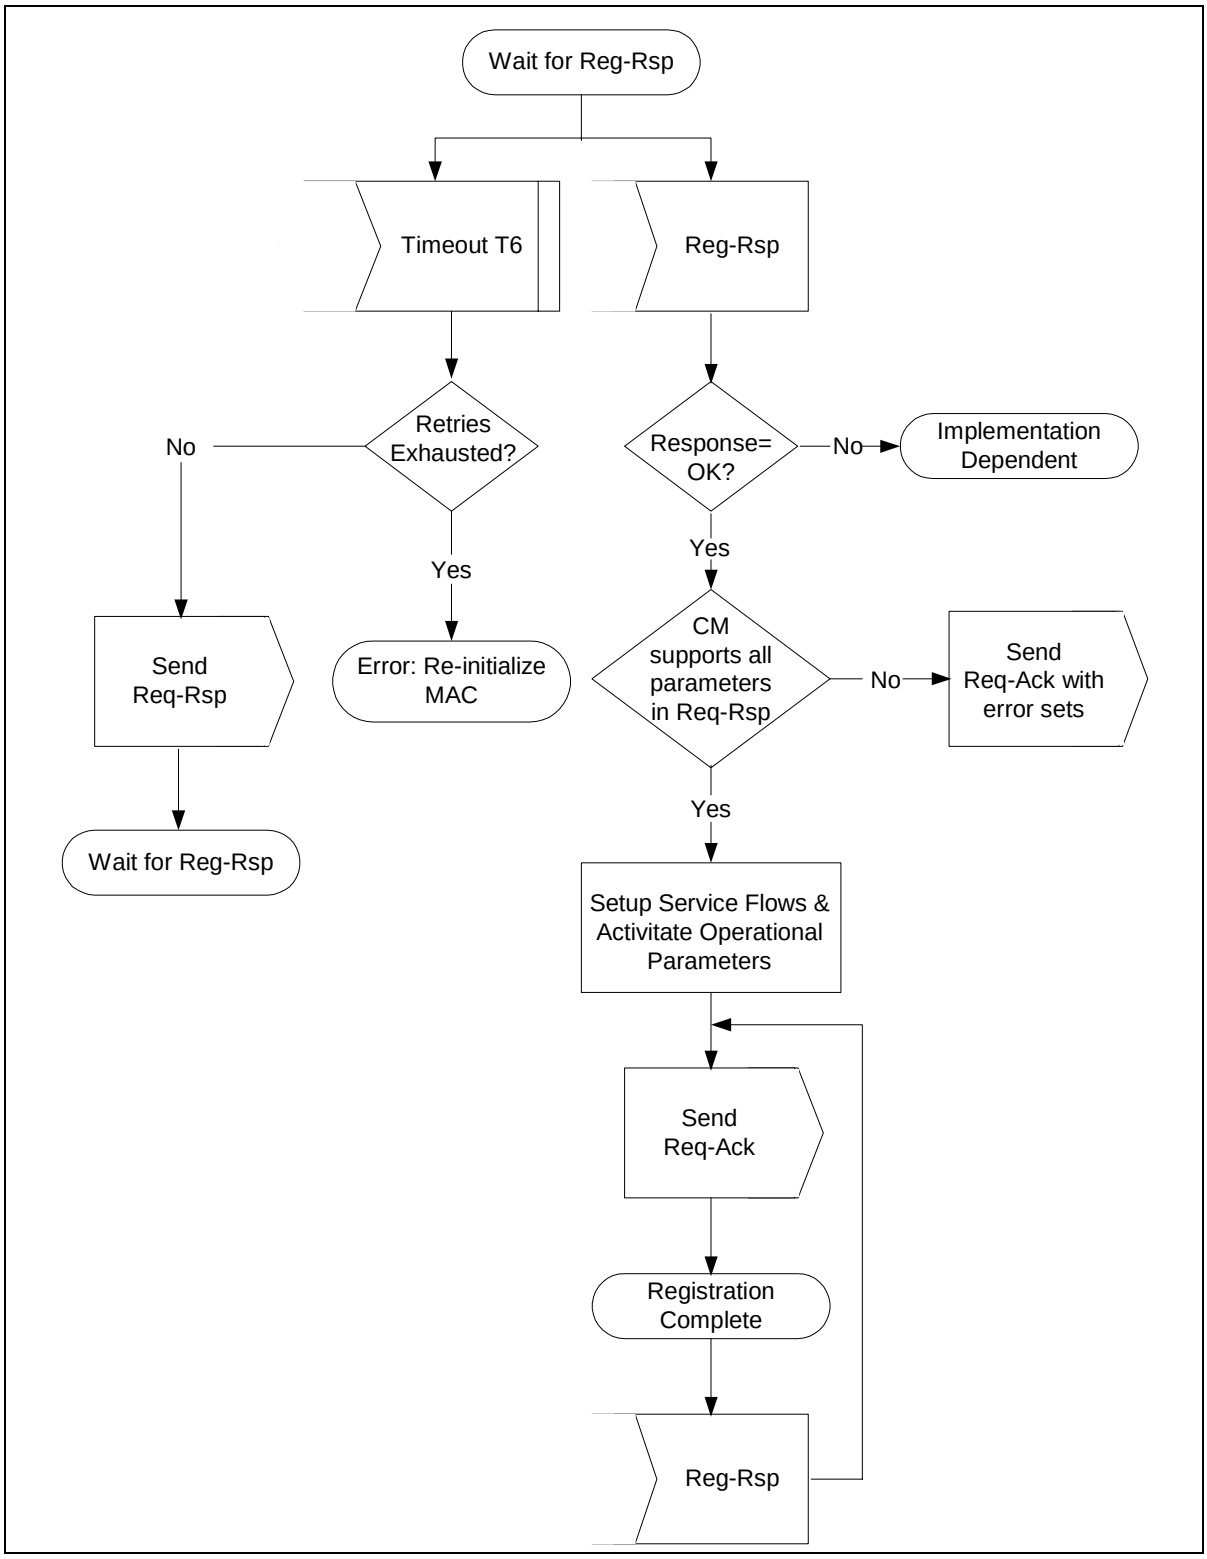


Figure 9-13. Wait for Registration Response — CM

The CMTS MUST perform the following operations to confirm the CM authorization (refer to Figure 9-14):

- Calculate a MIC per Appendix D.3.1 and compare it to the CMTS MIC included in the Registration Request. If the MIC is invalid, the CMTS MUST respond with an Authentication Failure.
- If present, 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 (refer to 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, 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.
- If the Registration Request contains DOCSIS 1.0 Class of Service encodings, verify the availability of the class(es) of service requested. If unable to provide the class(es) of service, the CMTS MUST respond with a Class of Service Failure and the appropriate Service Not Available response code(s). (refer to Appendix C.1.3.4)
- If the Registration Request contains Service Flow encodings, verify the availability of the Quality of Service requested in the provisioned Service Flow(s). If unable to provide the Service Flow(s), the CMTS MUST respond with either a reject-temporary or a reject-permanent (see Appendix C.4) and the appropriate Service Flow Response(s).
- If the Registration Request contains DOCSIS 1.0 Class of Service encodings and Service Flow encodings, the CMTS MUST respond with a Class of Service Failure and a Service Not Available response code set to 'reject-permanent' for all DOCSIS 1.0 Classes and Service Flows requested.
- 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' (refer to Section 6.3.8.1.1).
- Assign a Service Flow ID for each class of service supported.
- Reply to the modem in a Registration Response.
- If the Registration Request contains Service Flow encodings, and the REG-RESP was sent with a confirmation code of ok (0), the CMTS MUST wait for a Registration Acknowledgment as shown in Figure 9-14. If the Registration Request contains DOCSIS 1.0 Class of Service encodings, the CMTS MUST NOT wait for a Registration Acknowledgment.
- If timer T9 expires, the CMTS MUST both de-assign the temporary SID from that CM and make some provision for aging out that SID.

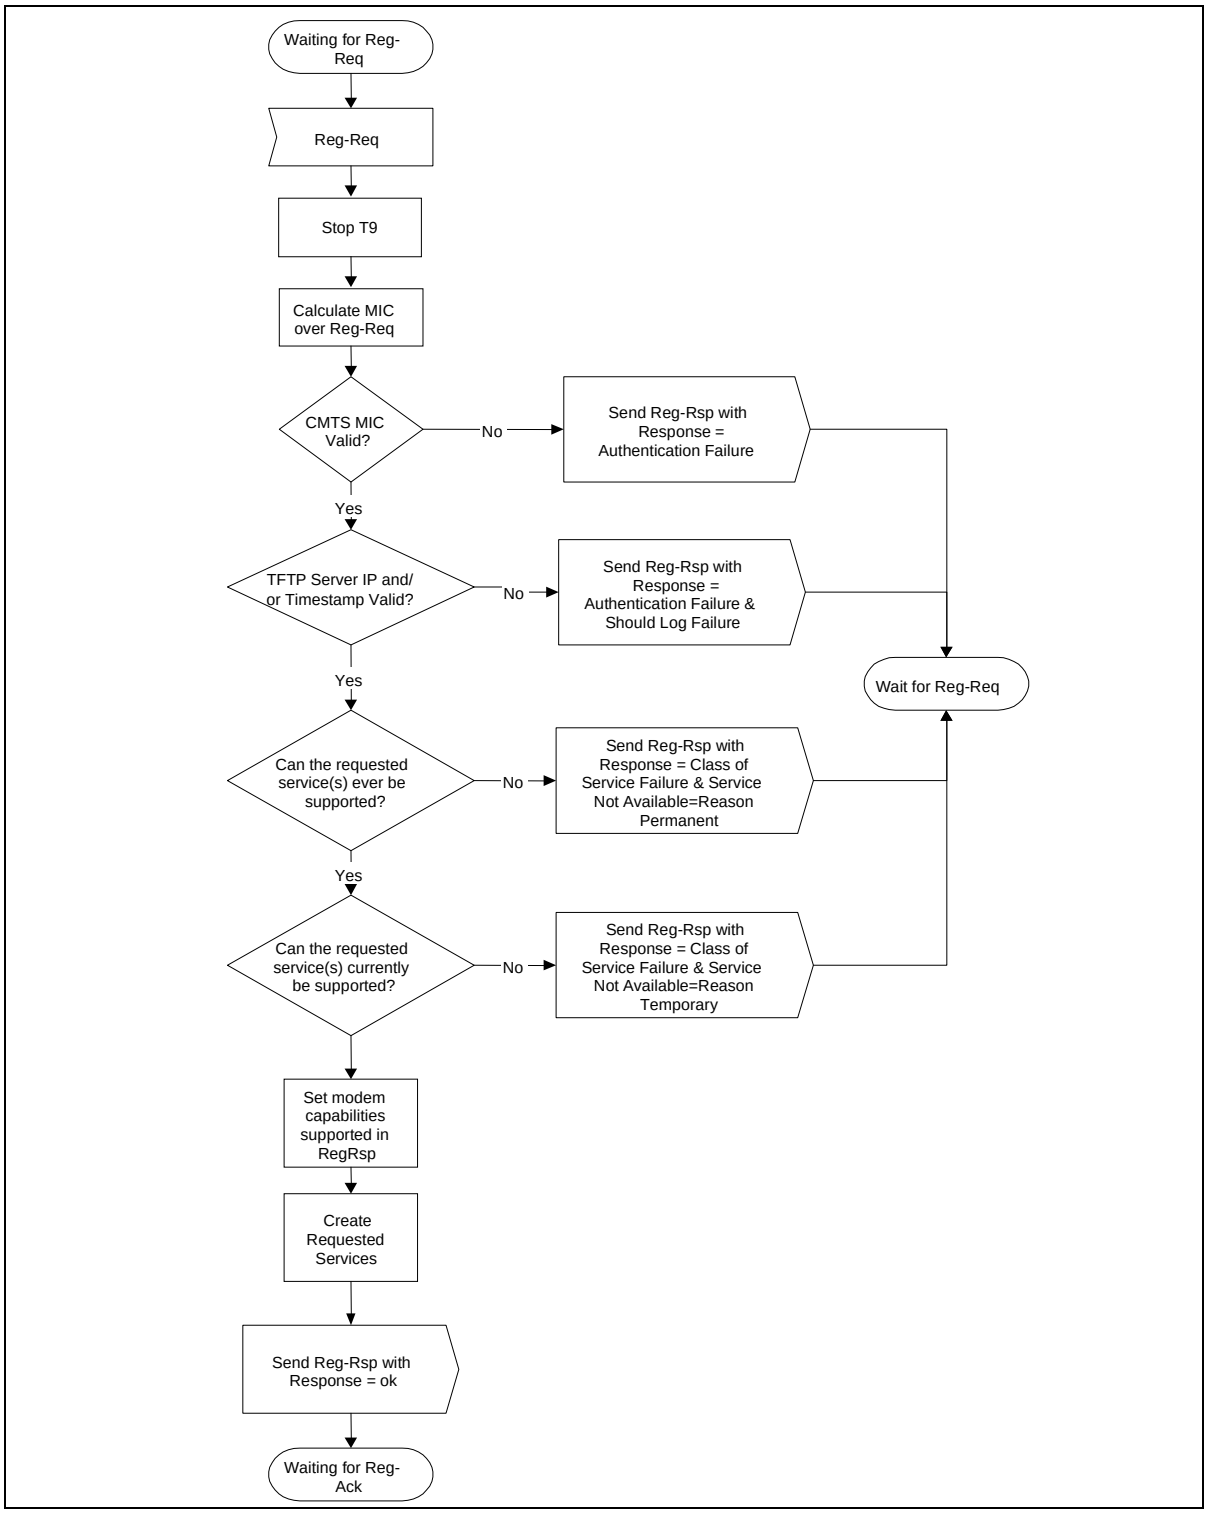


Figure 9-14. Registration — CMTS

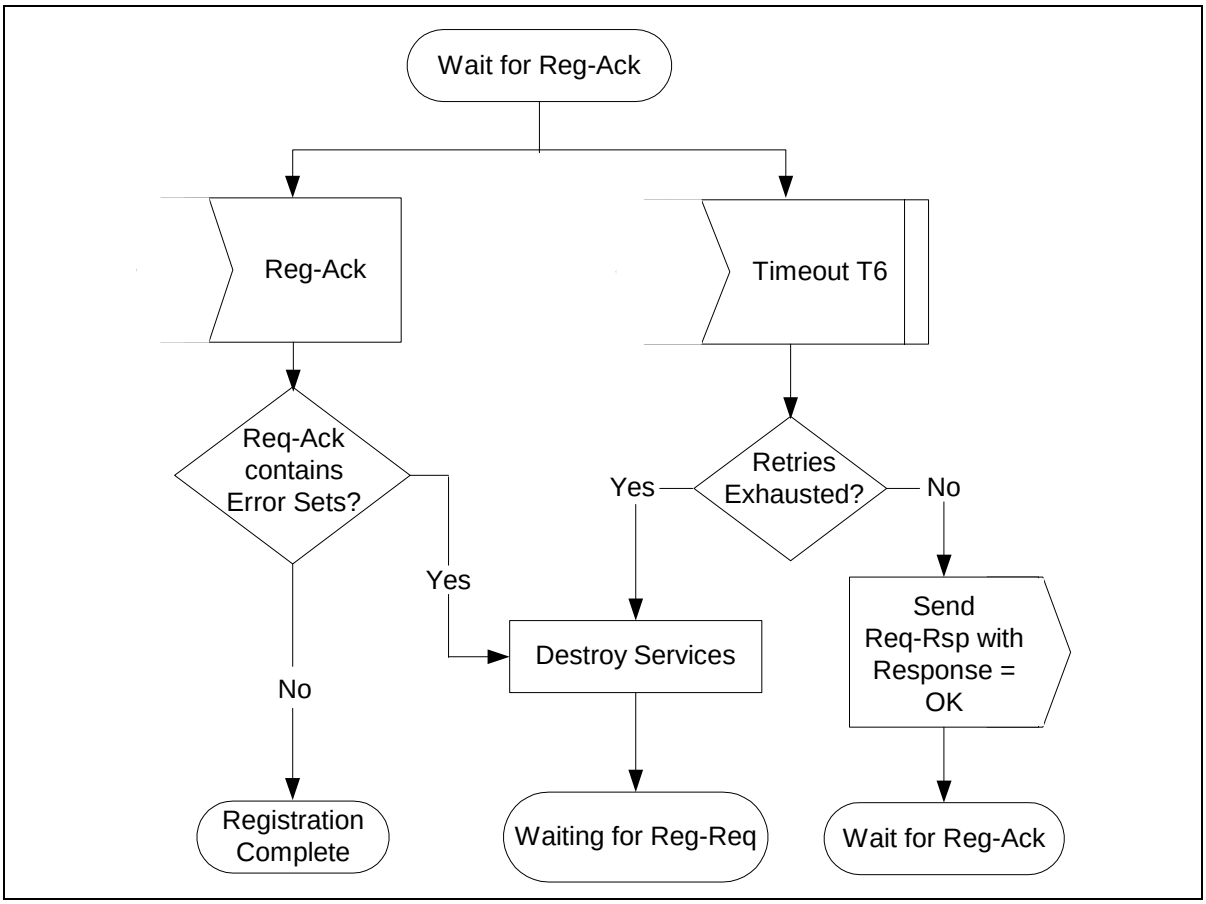


Figure 9-15. Registration Acknowledgment— CMTS

9.2.10 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]. A CM is provisioned to run Baseline Privacy if the Privacy Enable TLV (Appendix C.1.1.16) in the DOCSIS 1.1-style configuration file is explicitly/implicitly set to enable or the Baseline Privacy Configuration Setting (Appendix C.3.2) is contained in the DOCSIS 1.0-style configuration file as specified in Sections A.1.1 and C.2 of the BPI+ specification [DOCSIS8]. Note that the Secure Software Download is required regardless of whether the CM is provisioned to run Baseline Privacy or not as specified in Appendix D of the BPI+ specification [DOCSIS8]. If the CM has been provisioned to run Baseline Privacy, it MUST NOT forward traffic from any attached CPE device to the HFC network from the time registration completes until after the initialization of Baseline Privacy operations completes for its primary SID/SAID. The CMTS MAY stop forwarding data traffic to a CM that has been provisioned to run Baseline Privacy from the time registration completes until after the CM has successfully completed initializing Baseline Privacy.

9.2.11 Service IDs During CM Initialization

After completion of the Registration process (Section 9.2.9), the CM will have been assigned Service Flow IDs (SFIDs) to match its 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 6.3.8).

When assigning provisioned SFIDs on receiving a Registration Request, the CMTS may re-use the temporary SID, assigning it to one of the Service Flows 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.

9.2.12 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.

9.3 Standard Operation

9.3.1 Periodic Signal Level Adjustment

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.

Remote 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 9-16 and Figure 9-17. 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).

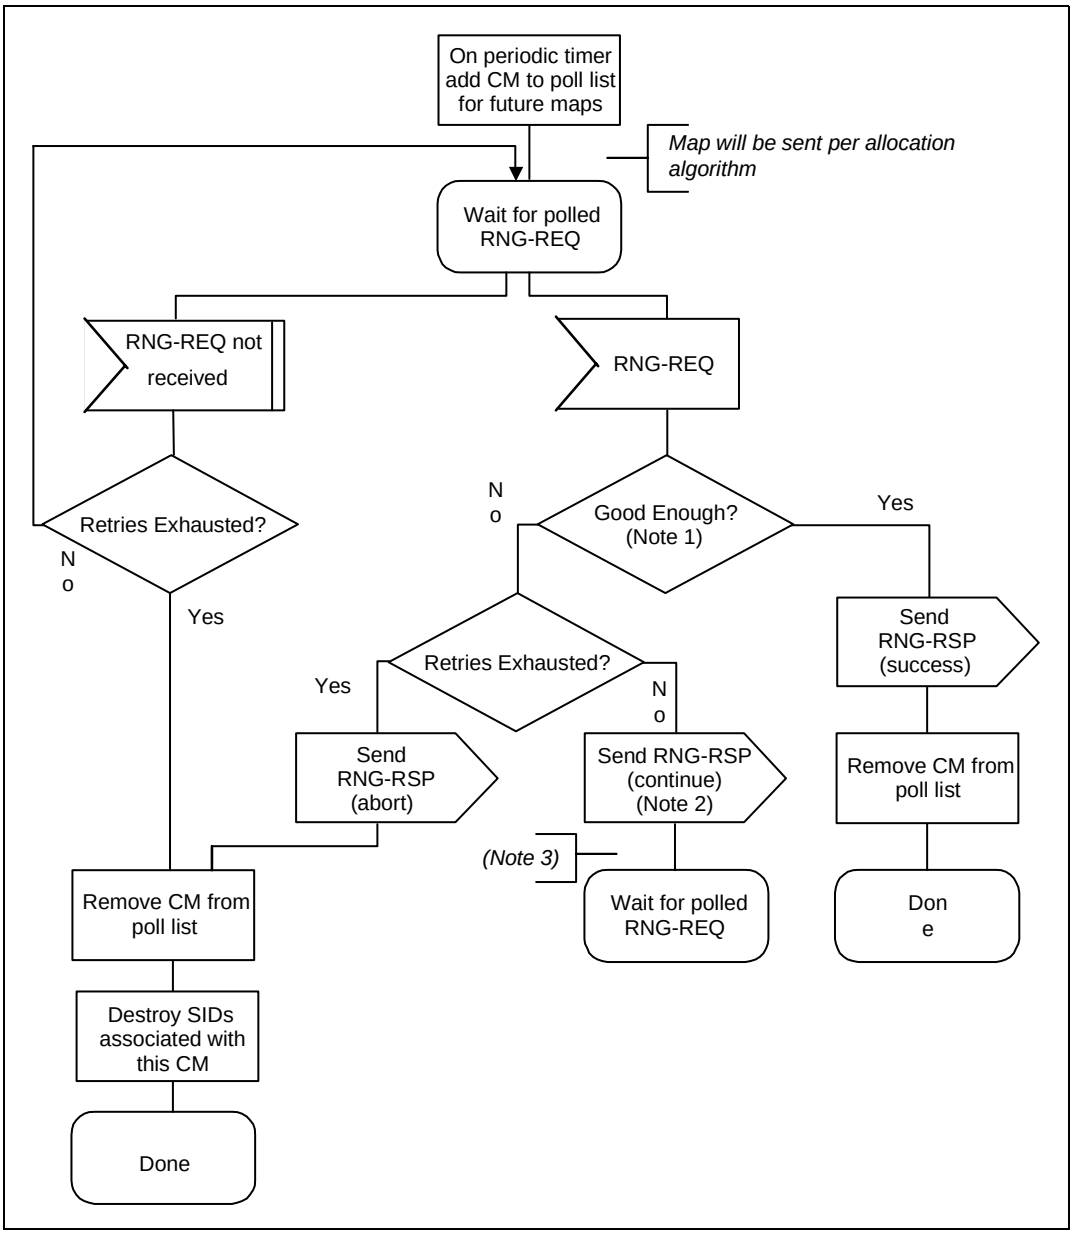


Figure 9-16. Periodic Ranging - CMTS

Notes to Figure 9-16:

1. Means pending-till-complete was zero and ranging is within the tolerable limits of the CMTS.
2. If pending-till-complete is nonzero, the CMTS MUST NOT send new Pre-Equalization Coefficients in the corresponding RNG-RSP. However, the CMTS MAY adjust other ranging parameters in the RNG-RSP message.
3. If the RNG-REQ pending-till-complete was nonzero, the CMTS MAY schedule additional station maintenance opportunities during the pending-till-complete period in order to adjust ranging parameters other than the equalizer.

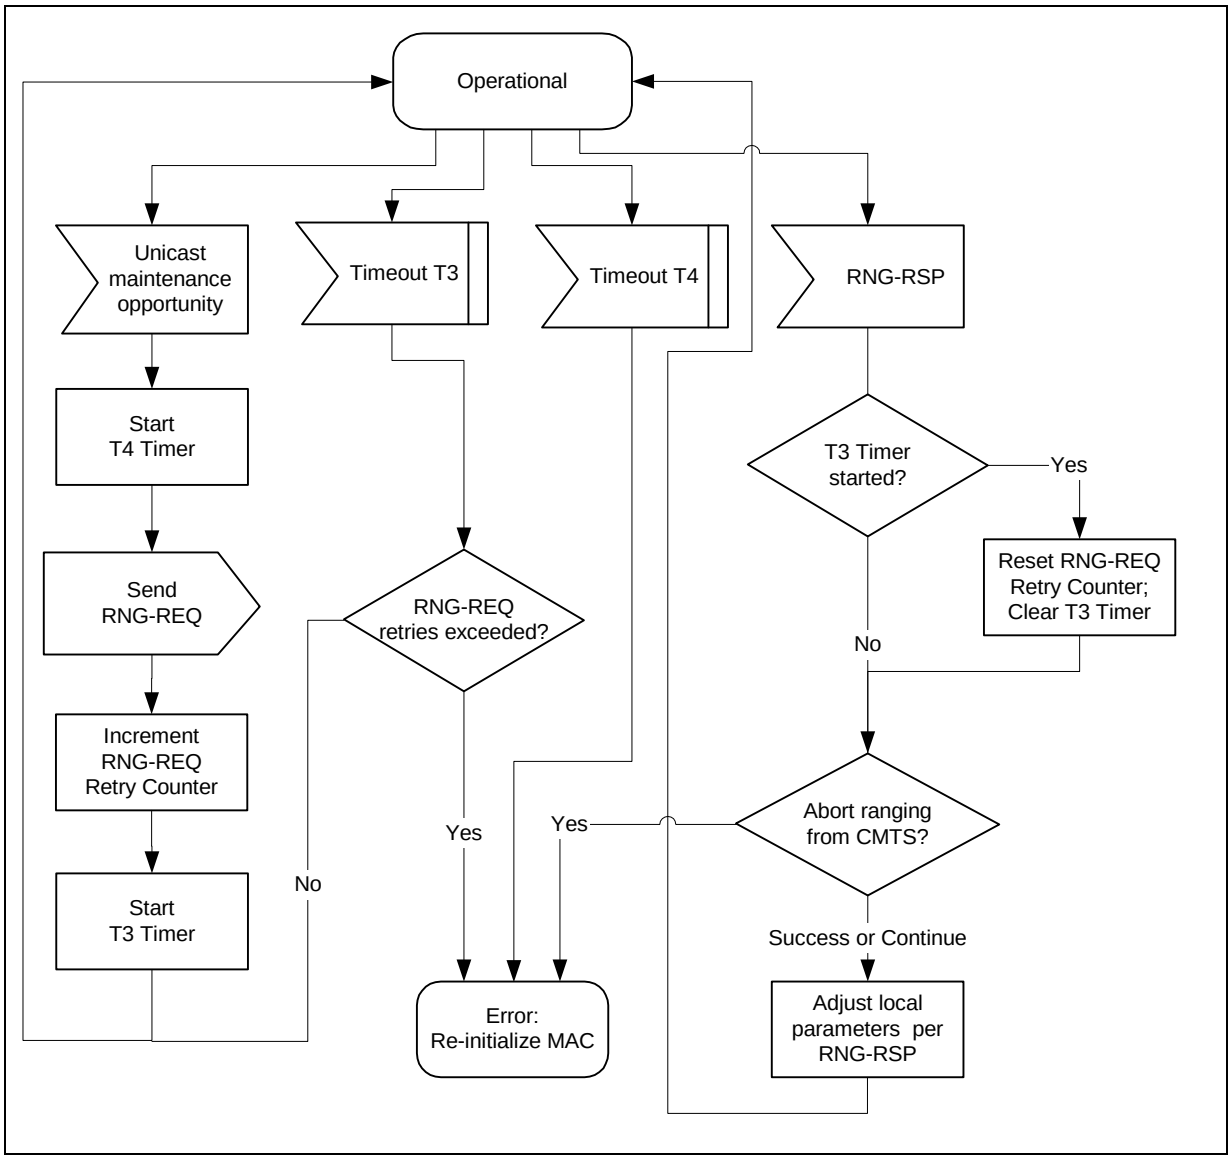


Figure 9-17. Periodic Ranging - CM View

9.3.2 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, and 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 7.1.5).

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.

9.3.3 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 9-18 shows the procedure that MUST be followed by the CMTS. Figure 9-19 shows the corresponding procedure at the CM.

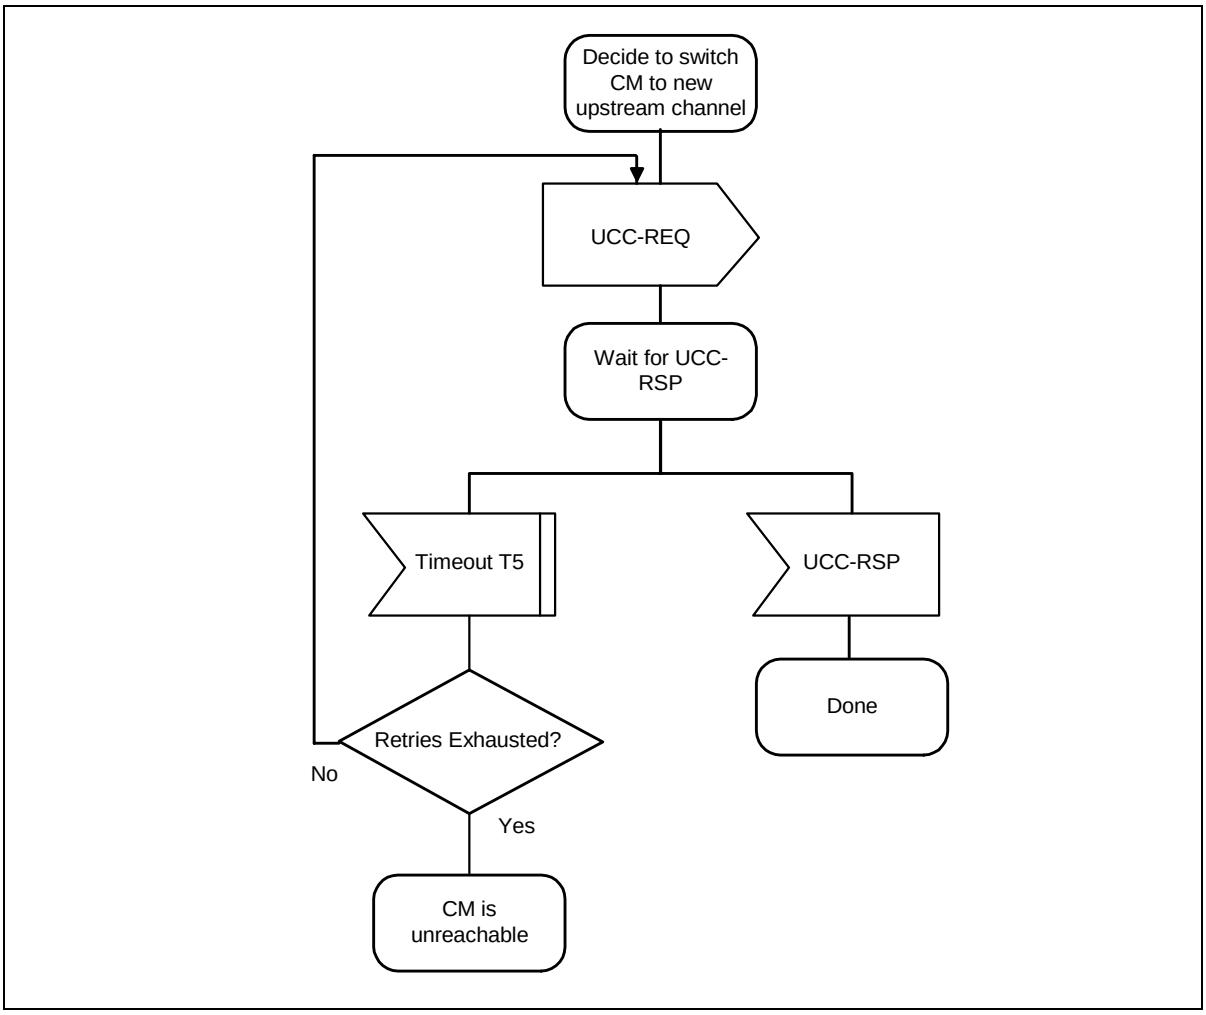


Figure 9-18. Changing Upstream Channels: CMTS View

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.

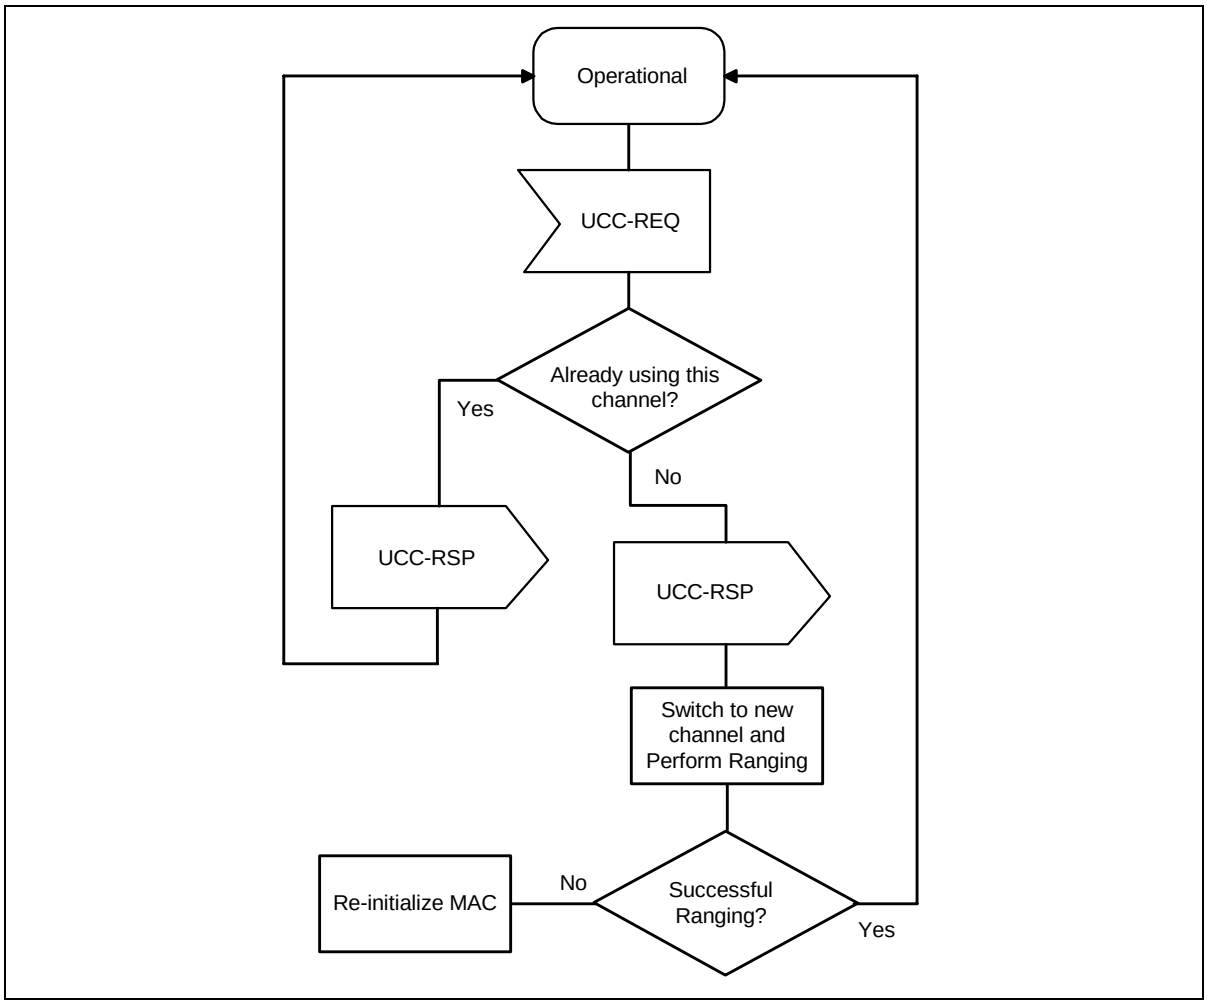


Figure 9-19. Changing Upstream Channels: CM View

Upon synchronizing with the new upstream channel, the CM MUST perform initial maintenance on the new upstream 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), then the CM MAY use cached ranging information and omit ranging.

The CM SHOULD cache UCD information from multiple upstream channels to eliminate waiting for a UCD corresponding to the new upstream channel.

The CM MUST NOT perform re-registration, since its provisioning and MAC domain remain valid on the new channel.

9.4 Dynamic Service

Service Flows may be created, changed or deleted. This is accomplished through a series of MAC management messages referred to as Dynamic Service Addition (DSA), Dynamic Service Change (DSC) and Dynamic Service Deletion (DSD). The DSA messages create a new Service Flow. The DSC messages change an existing

Service Flow. The DSD messages delete a single existing Upstream and/or a single existing Downstream Service Flow. This is illustrated in Figure 9-20.

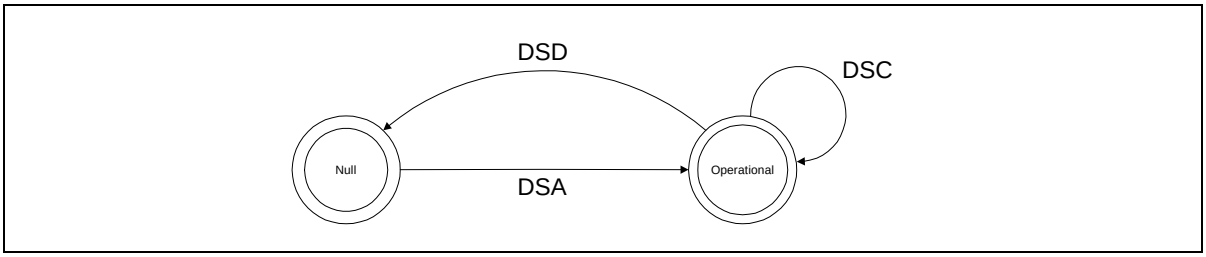


Figure 9-20. Dynamic Service Flow Overview

The Null state implies that no Service Flow exists that matches the SFID and/or TransactionID in a message. Once the Service Flow exists, it is operational and has an assigned SFID. In steady state operation, a Service Flow resides in a Nominal state. When Dynamic Service messaging is occurring, the Service Flow may transition through other states, but remains operational. Since multiple Service Flows may exist, there may be multiple state machines active, one for every Service Flow. Dynamic Service messages only affect those state machines that match the SFID and/or Transaction ID. If privacy is enabled, both the CM and CMTS MUST verify the HMAC digest on all dynamic service messages before processing them, and discard any messages that fail.

Service Flows created at registration time effectively enter the SF_operational state without a DSA transaction.

TransactionIDs are unique per transaction and are selected by the initiating device (CM or CMTS). To help prevent ambiguity and provide simple checking, the TransactionID number space is split between the CM and CMTS. The CM MUST select its TransactionIDs from the first half of the number space (0x0000 to 0x7FFF). The CMTS MUST select its TransactionIDs from the second half of the number space (0x8000 to 0xFFFF).

Each dynamic service message sequence is a unique transaction with an associated unique transaction identifier. The DSA/DSC transactions consist of a request/response/acknowledge sequence. The DSD transactions consist of a request/response sequence. The response messages MUST contain a confirmation code of okay unless some exception condition was detected. The acknowledge messages MUST include the confirmation code in the response unless a new exception condition arises. A more detailed state diagram, including transition states, is shown below. The detailed actions for each transaction will be given in the following sections.

9.4.1 Dynamic Service Flow State Transitions

The Dynamic Service Flow State Transition Diagram is the top-level state diagram and controls the general Service Flow state. As needed, it creates transactions, each represented by a Transaction state transition diagram, to provide the DSA, DSC, and DSD signaling. Each Transaction state transition diagram only communicates with the parent Dynamic Service Flow State Transition Diagram. The top-level state transition diagram filters Dynamic Service messages and passes them to the appropriate transaction based on Service Flow Identifier (SFID), Service Flow Reference number, and TransactionID.

If a single Dynamic Service message affects a pair of service flows, a single transaction is initiated which communicates with both parent Dynamic Service Flow State Transition Diagrams. In this case, both service flows MUST remain locked in the same state until they receive the DSx Succeeded or DSx Failed input from the DSx Transaction State Transition Diagram. During that “lock interval”, if a message is received which refers to only one of the two service flows, it MUST be treated as if it refers to both service flows, so that both service flows stay in the same state. If a DSD-REQ message is received during the lock interval which refers to only one of the two service flows, the device MUST handle the event normally, by sending the SF Delete-Remote to the

ongoing DSx Transaction and by initiating a DSD-Remote transaction, and in addition, it MUST initiate a DSD-Local transaction to delete the second service flow of the locked pair.

If a DSC request is received which refers to two service flows locked in different transactions, and they are in different states, the device MUST reject the request without affecting the ongoing transactions.

There are six different types of transactions: locally initiated or remotely initiated for each of the DSA, DSC and DSD messages. Most transactions have three basic states: pending, holding and deleting. The pending state is typically entered after creation and is where the transaction is waiting for a reply. The holding state is typically entered once the reply is received. the purpose of this state is to allow for retransmissions in case of a lost message, even though the local entity has perceived that the transaction has completed. The deleting state is only entered if the Service Flow is being deleted while a transaction is being processed.

The flow diagrams provide a detailed representation of each of the states in the Transaction state transition diagrams. All valid transitions are shown. Any inputs not shown should be handled as a severe error condition.

With one exception, these state diagrams apply equally to the CMTS and CM. In the Dynamic Service Flow Changing-Local state, there is a subtle difference in the CM and CMTS behaviors. This is called out in the state transition and detailed flow diagrams.

Note: The 'Num Xacts' variable in the Dynamic Service Flow State Transition Diagram is incremented every time the top-level state diagram creates a transaction and is decremented every time a transaction terminates. A Dynamic Service Flow MUST NOT return to the Null state until it's deleted and all transactions have terminated.

The inputs for the state diagrams are identified below.

Dynamic Service Flow State Transition Diagram inputs from unspecified local, higher-level entities:

- Add
- Change
- Delete

Dynamic Service Flow State Transition Diagram inputs from DSx Transaction State Transition diagrams:

- DSA Succeeded
- DSA Failed
- DSA ACK Lost
- DSA Erred
- DSA Ended

- DSC Succeeded
- DSC Failed
- DSC ACK Lost
- DSC Erred
- DSC Ended

- DSD Succeeded
- DSD Erred

- DSD Ended

DSx Transaction State Transition diagram inputs from the Dynamic Service Flow State Transition Diagram:

- SF Add
- SF Change
- SF Delete

- SF Abort Add
- SF Change-Remote
- SF Delete-Local
- SF Delete-Remote

- SF DSA-ACK Lost
- SF-DSC-REQ Lost
- SF-DSC-ACK Lost
- SF DSD-REQ Lost

- SF Changed
- SF Deleted

The creation of DSx Transactions by the Dynamic Service Flow State Transition Diagram is indicated by the notation

DSx-[Local | Remote] (initial_input)

where initial_input may be SF Add, DSA-REQ, SF Change, DSC-REQ, SF Delete, or DSD-REQ depending on the transaction type and initiator.

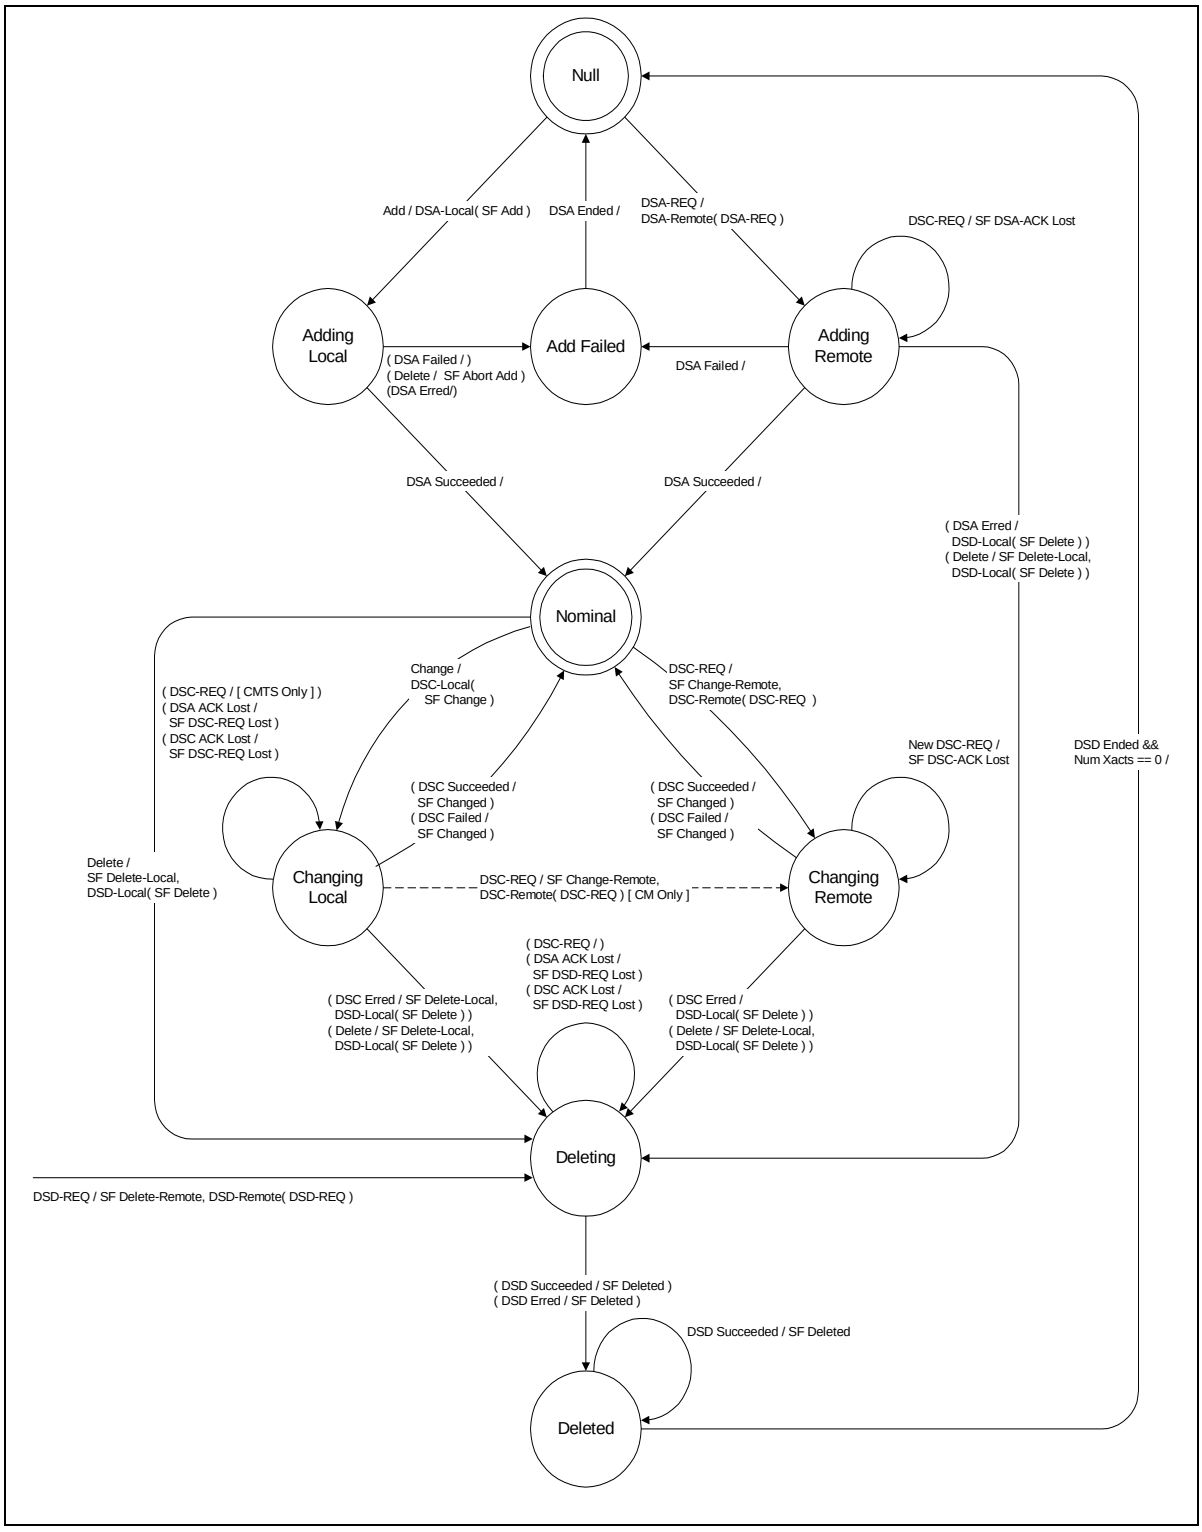


Figure 9-21. Dynamic Service Flow State Transition Diagram

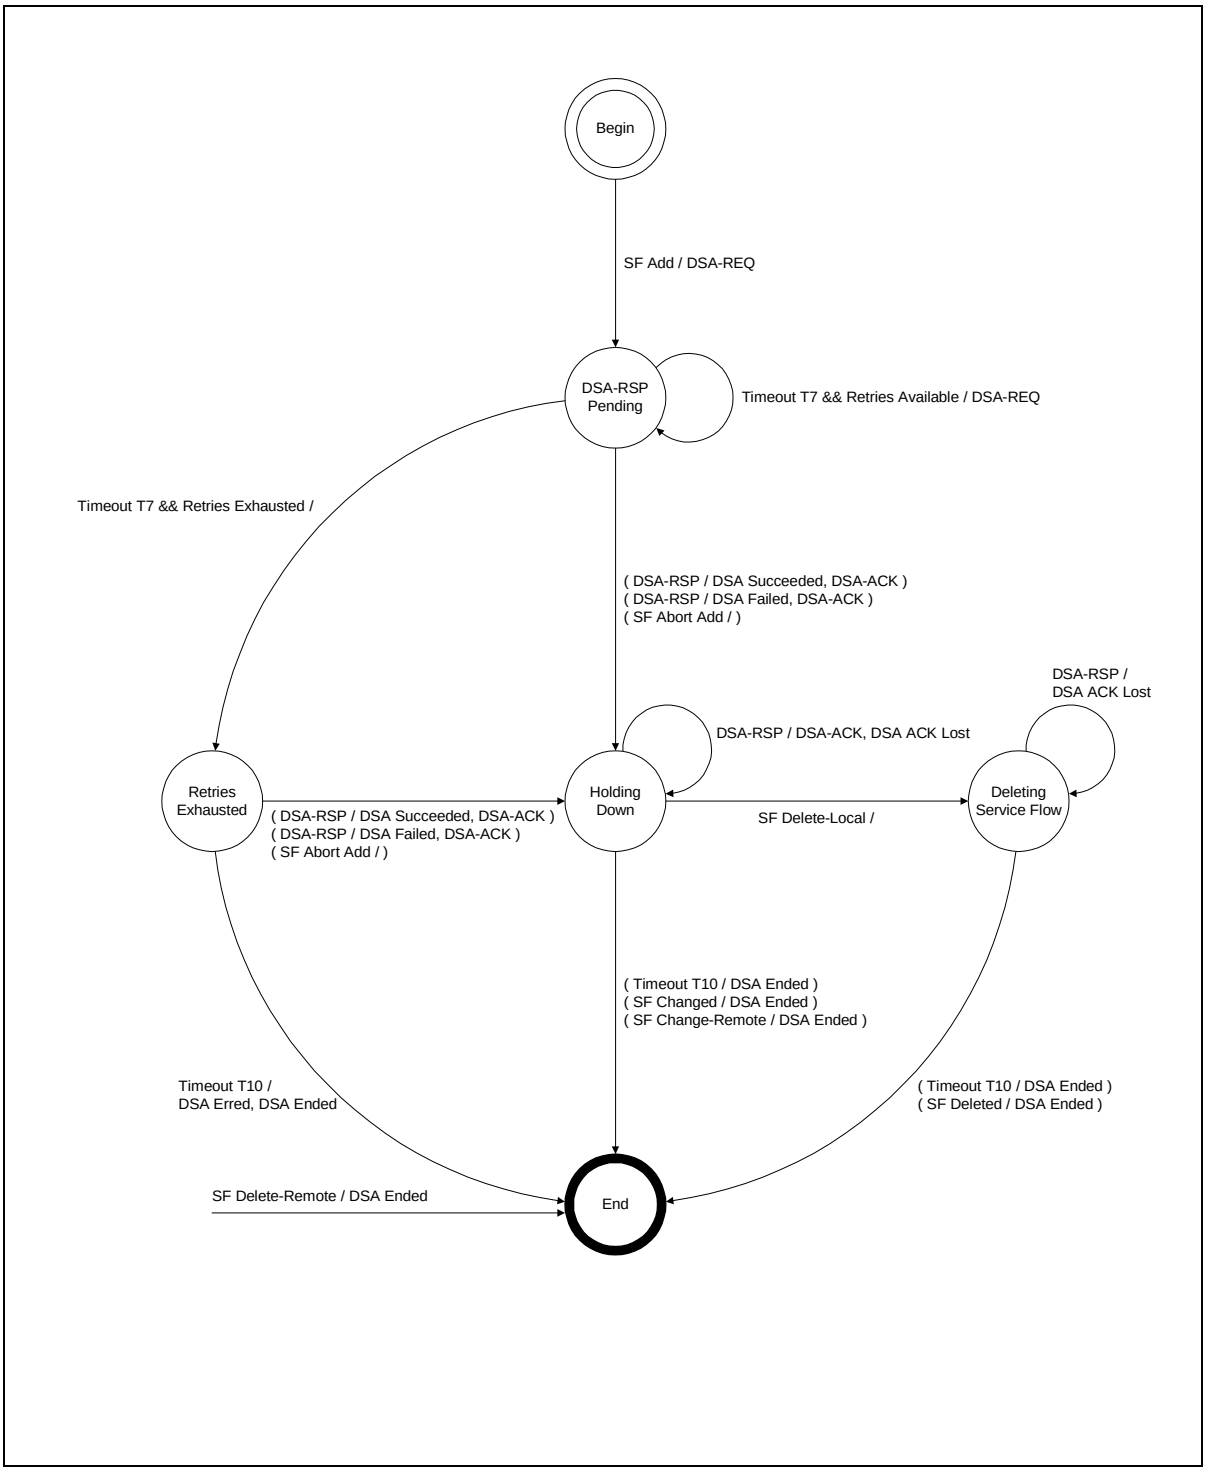


Figure 9-22. DSA—Locally Initiated Transaction State Transition Diagram

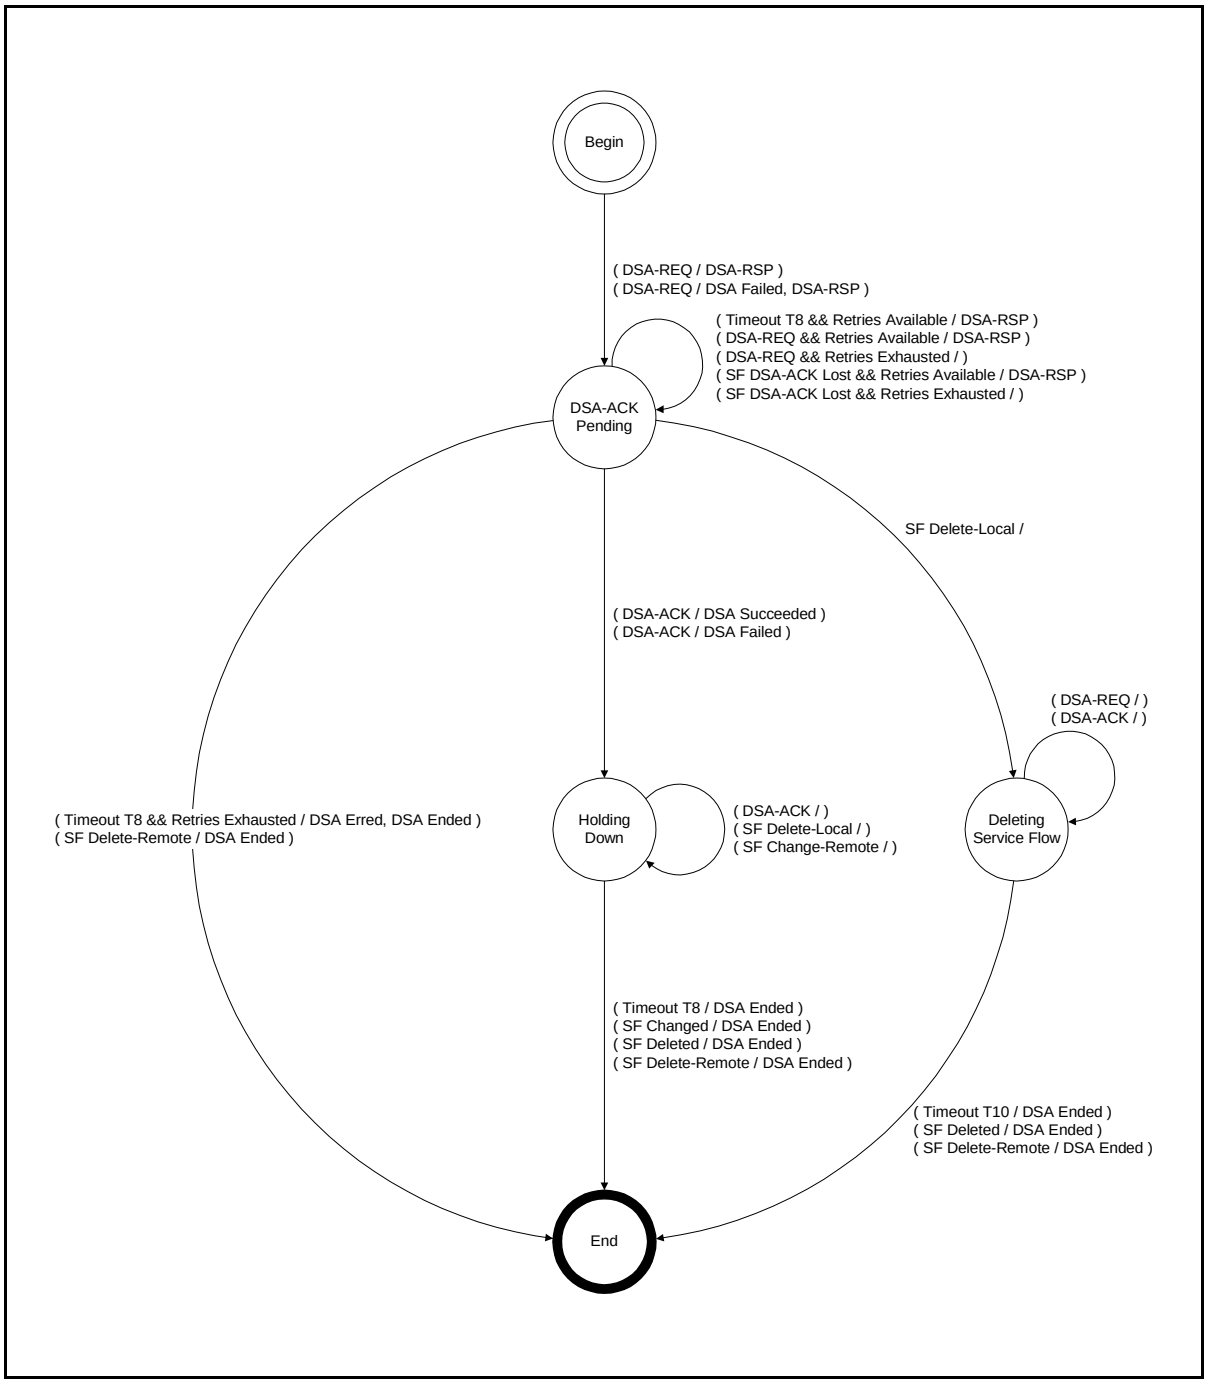


Figure 9-23. DSA—Remotely Initiated Transaction State Transition Diagram

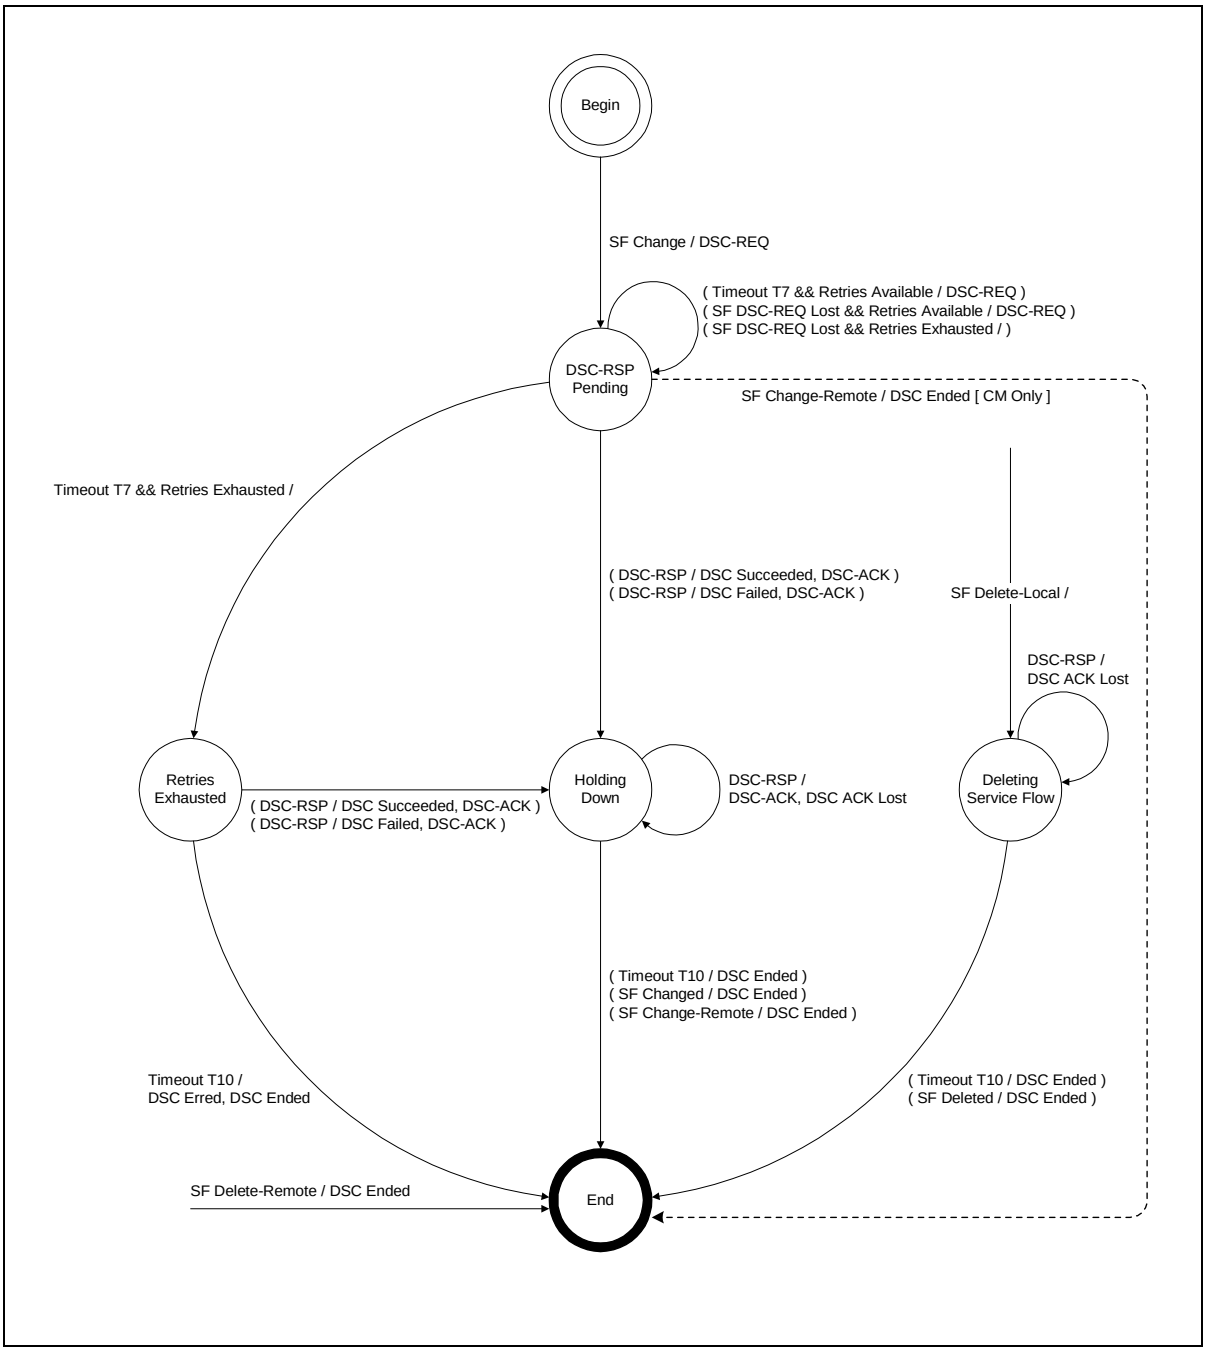


Figure 9-24. DSC—Locally Initiated Transaction State Transition Diagram

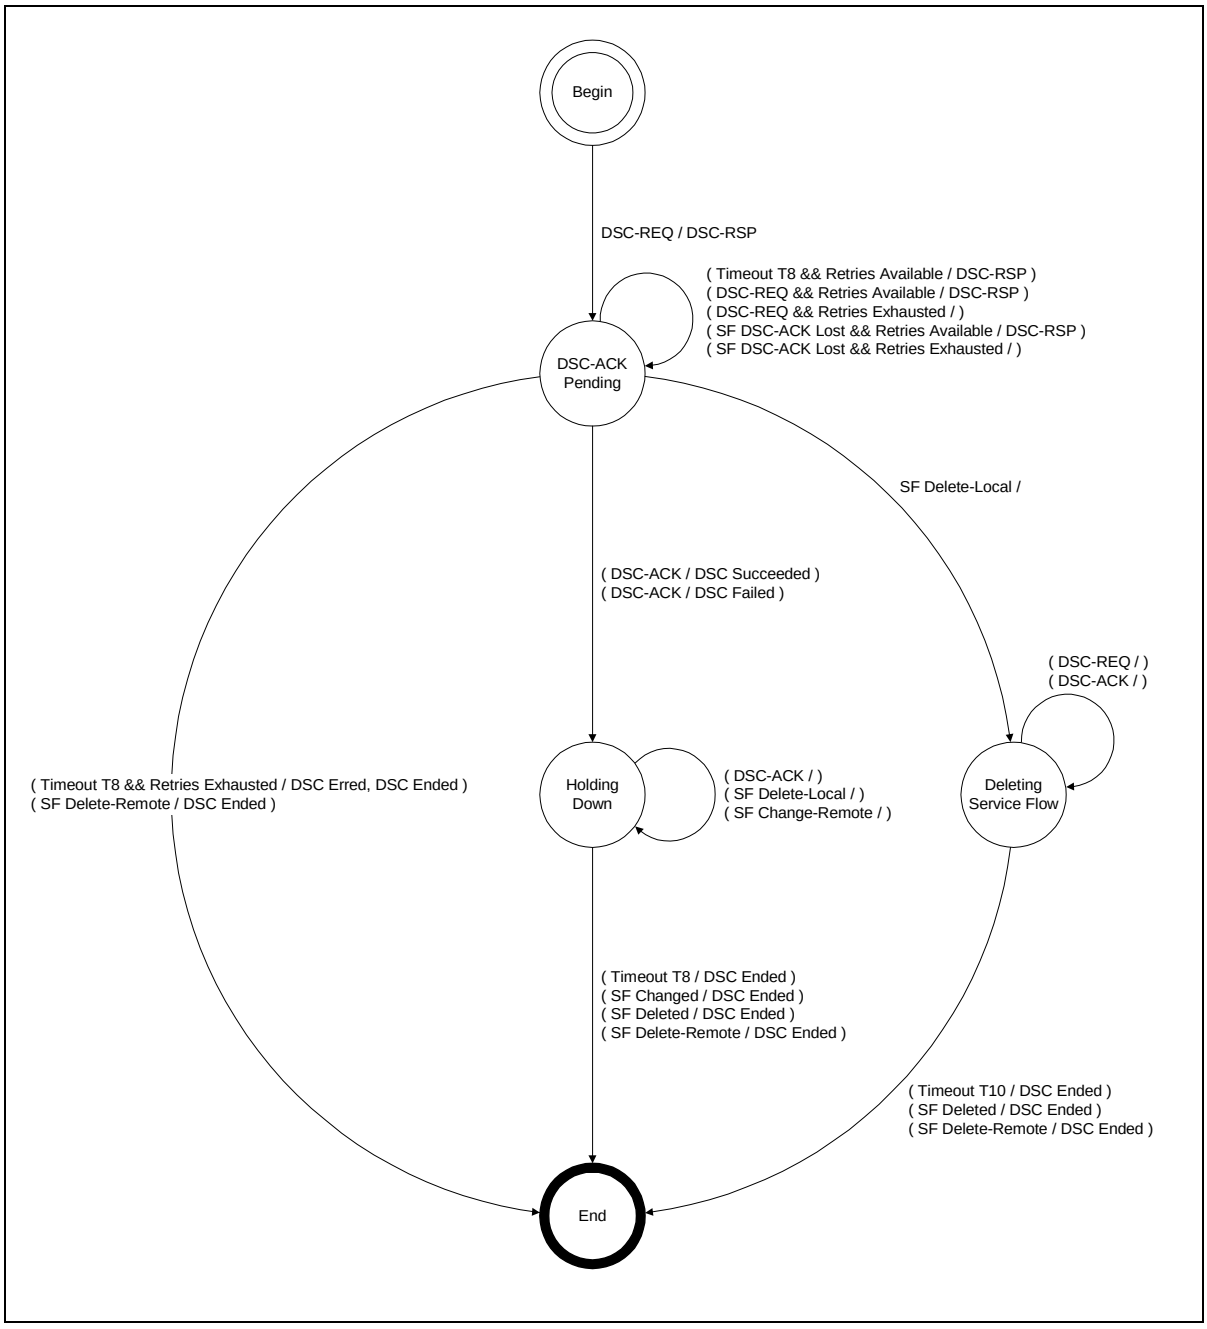


Figure 9-25. DSC—Remotely Initiated Transaction State Transition Diagram

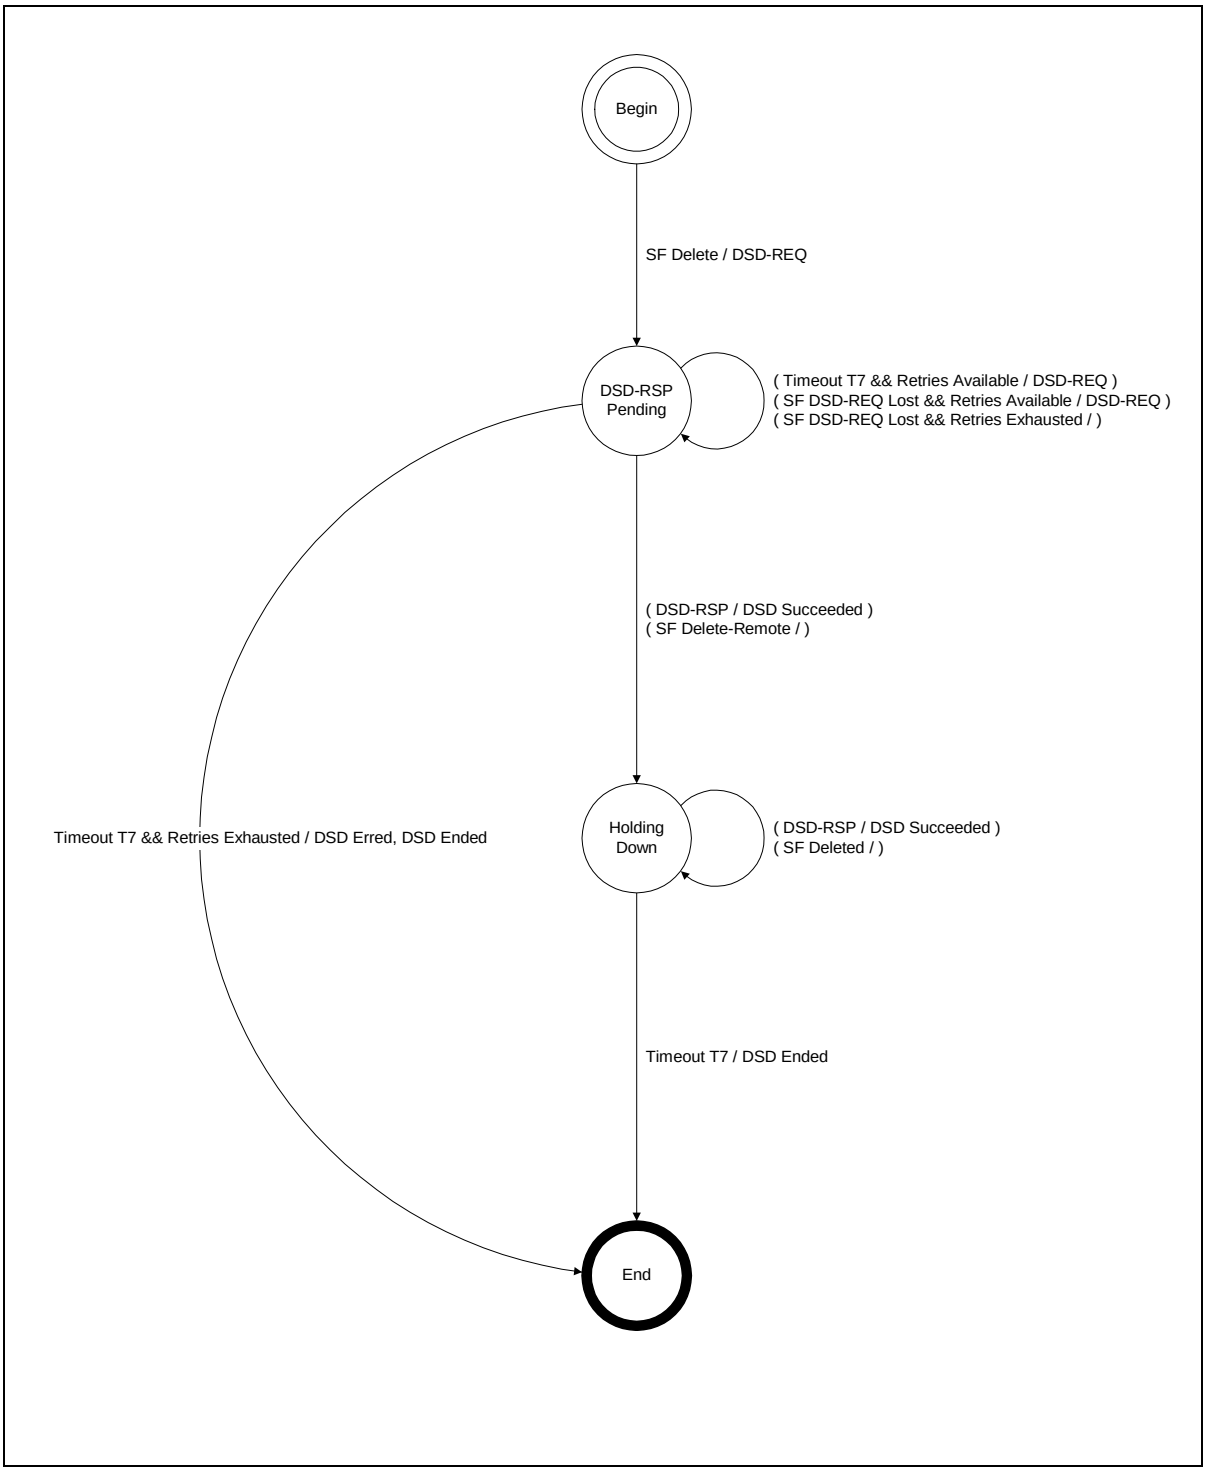


Figure 9-26. DSD—Locally Initiated Transaction State Transition Diagram

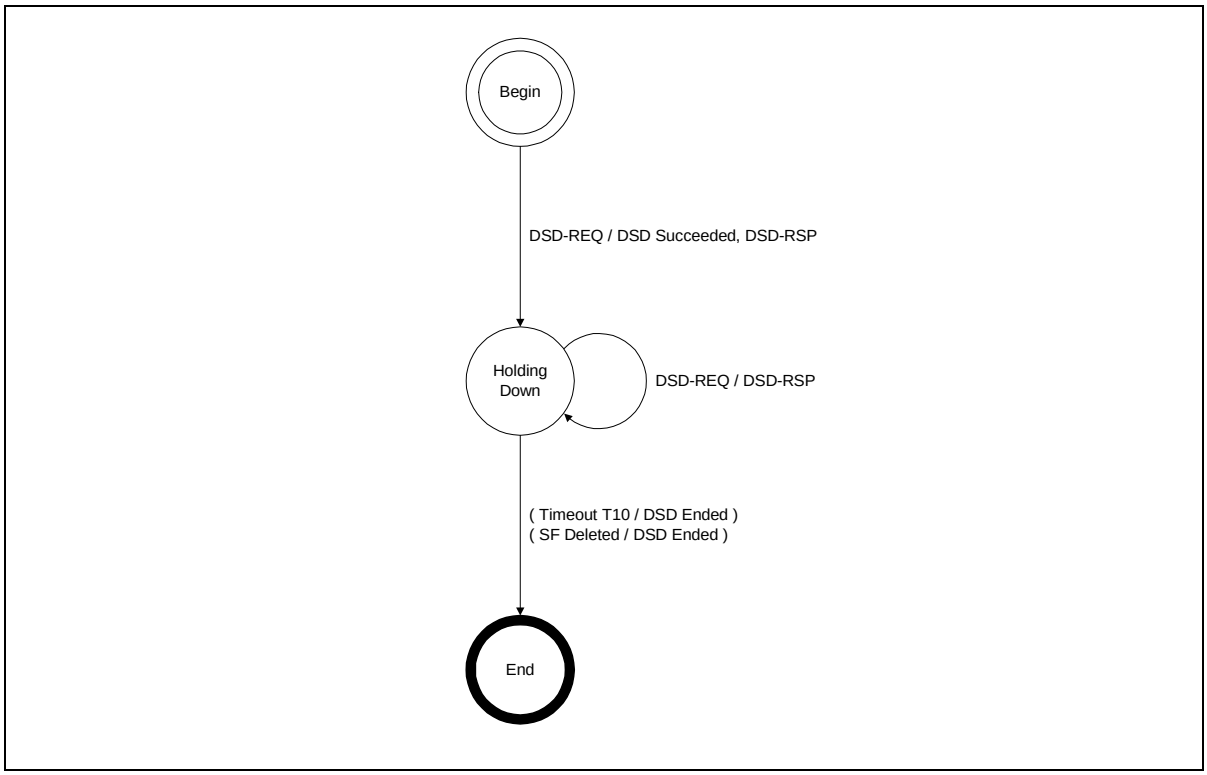


Figure 9-27. Dynamic Deletion (DSD) - Remotely Initiated Transaction State Transition Diagram

9.4.2 Dynamic Service Addition

9.4.2.1 CM Initiated Dynamic Service Addition

A CM wishing to create an upstream and/or a downstream Service Flow sends a request to the CMTS using a dynamic service addition request message (DSA-REQ). The CMTS checks the CM's authorization for the requested service(s) and whether the QoS requirements can be supported and generates an appropriate response using a dynamic service addition response message (DSA-RSP). The CM concludes the transaction with an acknowledgment message (DSA-ACK).

In order to facilitate a common admission response, an upstream and a downstream Service Flow can be included in a single DSA-REQ. Both Service Flows are either accepted or rejected together.

CM		CMTS
New Service Flow(s) needed		
Check if resources are available		
Send DSA-REQ	**---DSA-REQ-->**	**Receive DSA-REQ**
		Check if CM authorized for Service(s)[1]
		Check Service Flow(s) QoS can be supported
		Create SFID(s)
		If upstream AdmittedQoSParamSet is non-null, Create SID
		If upstream ActiveQoSParamSet is non-null, Enable reception of data on new upstream Service Flow
Receive DSA-RSP	**<--DSA-RSP---**	**Send DSA-RSP**
If ActiveQoSParamSet is non-null, Enable transmission and/or reception of data on new Service Flow(s)		
Send DSA-ACK	**---DSA-ACK-->**	**Receive DSA-ACK**
		If downstream ActiveQoSParamSet is non-null, Enable transmission of data on new downstream Service Flow

Figure 9-28. Dynamic Service Addition Initiated from CM

Notes to Figure 9-28:

1. Authorization can happen prior to the DSA-REQ being received by the CMTS. The details of CMTS signalling to anticipate a DSA-REQ are beyond the scope of this specification.

9.4.2.2 CMTS Initiated Dynamic Service Addition

A CMTS wishing to establish an upstream and/or a downstream dynamic Service Flow(s) with a CM performs the following operations. The CMTS checks the authorization of the destination CM for the requested class of service and whether the QoS requirements can be supported. If the service can be supported the CMTS generates new SFID(s) with the required class of service and informs the CM using a dynamic service addition request message (DSA-REQ). If the CM checks that it can support the service and responds using a dynamic service addition response message (DSA-RSP). The transaction completes with the CMTS sending the acknowledge message (DSA-ACK).

CM		CMTS
		New Service Flow(s) required for CM
		Check CM authorized for Service(s)
		Check Service Flow(s) QoS can be supported
		Create SFID(s)
		If upstream AdmittedQoSParamSet is non-null, Create SID
		If upstream ActiveQoSParamSet is non-null, Enable reception of data on new upstream Service Flow
Receive DSA-REQ	**<--DSA-REQ---**	**Send DSA-REQ**
Confirm CM can support Service Flow(s)		
Add Downstream SFID (if present)		
Enable reception on any new downstream Service Flow		
Send DSA-RSP	**---DSA-RSP-->**	**Receive DSA-RSP**
		Enable transmission & reception of data on new Service Flow(s)
Receive DSA-ACK	**<--DSA-ACK---**	**Send DSA-ACK**
Enable transmission on new upstream Service Flow		

Figure 9-29. Dynamic Service Addition Initiated from CMTS

9.4.2.3 Dynamic Service Addition State Transition Diagrams

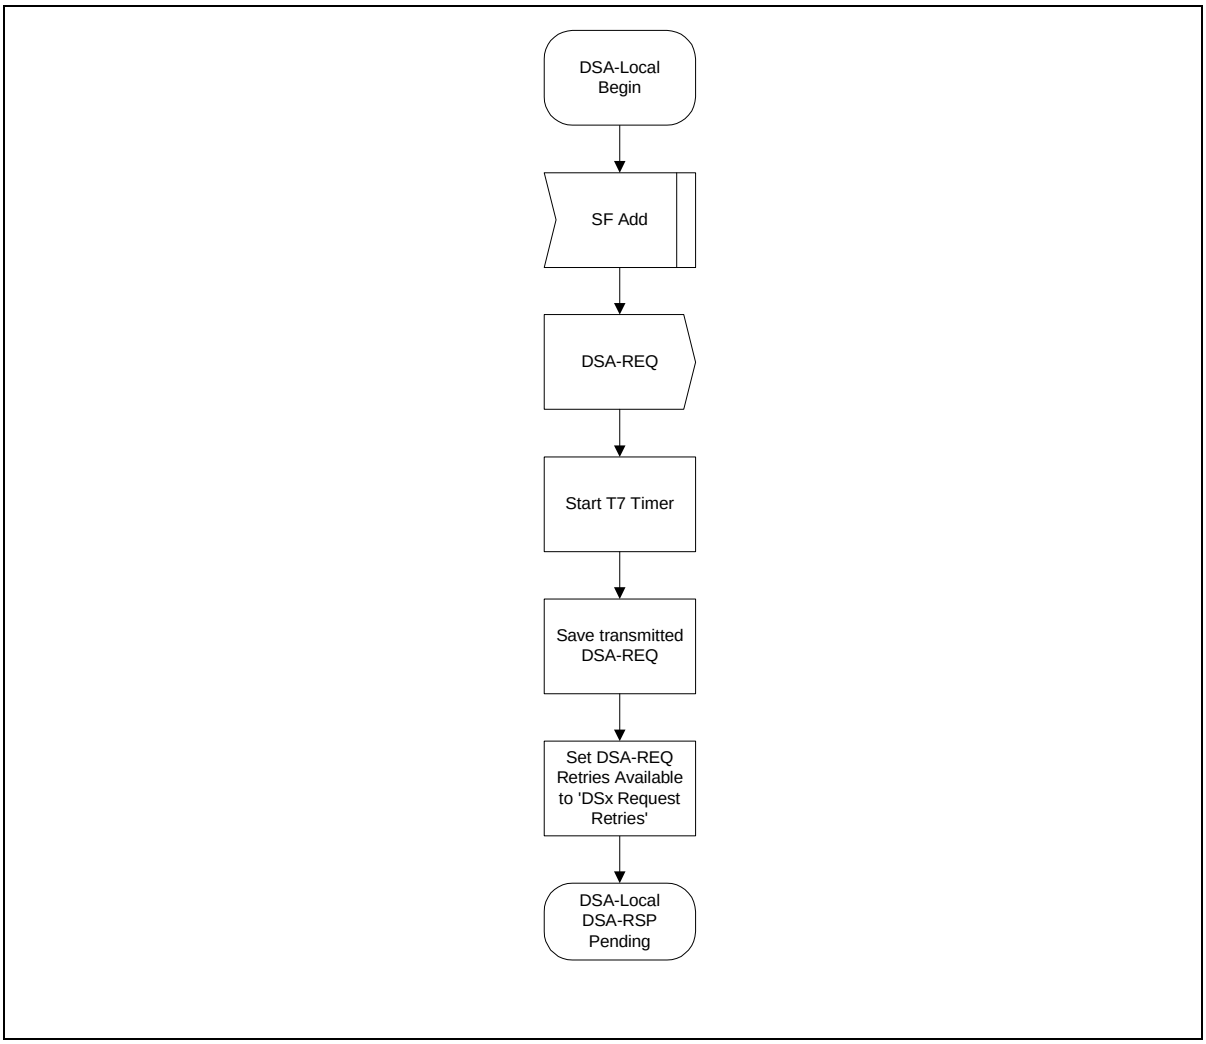


Figure 9-30. DSA—Locally Initiated Transaction Begin State Flow Diagram

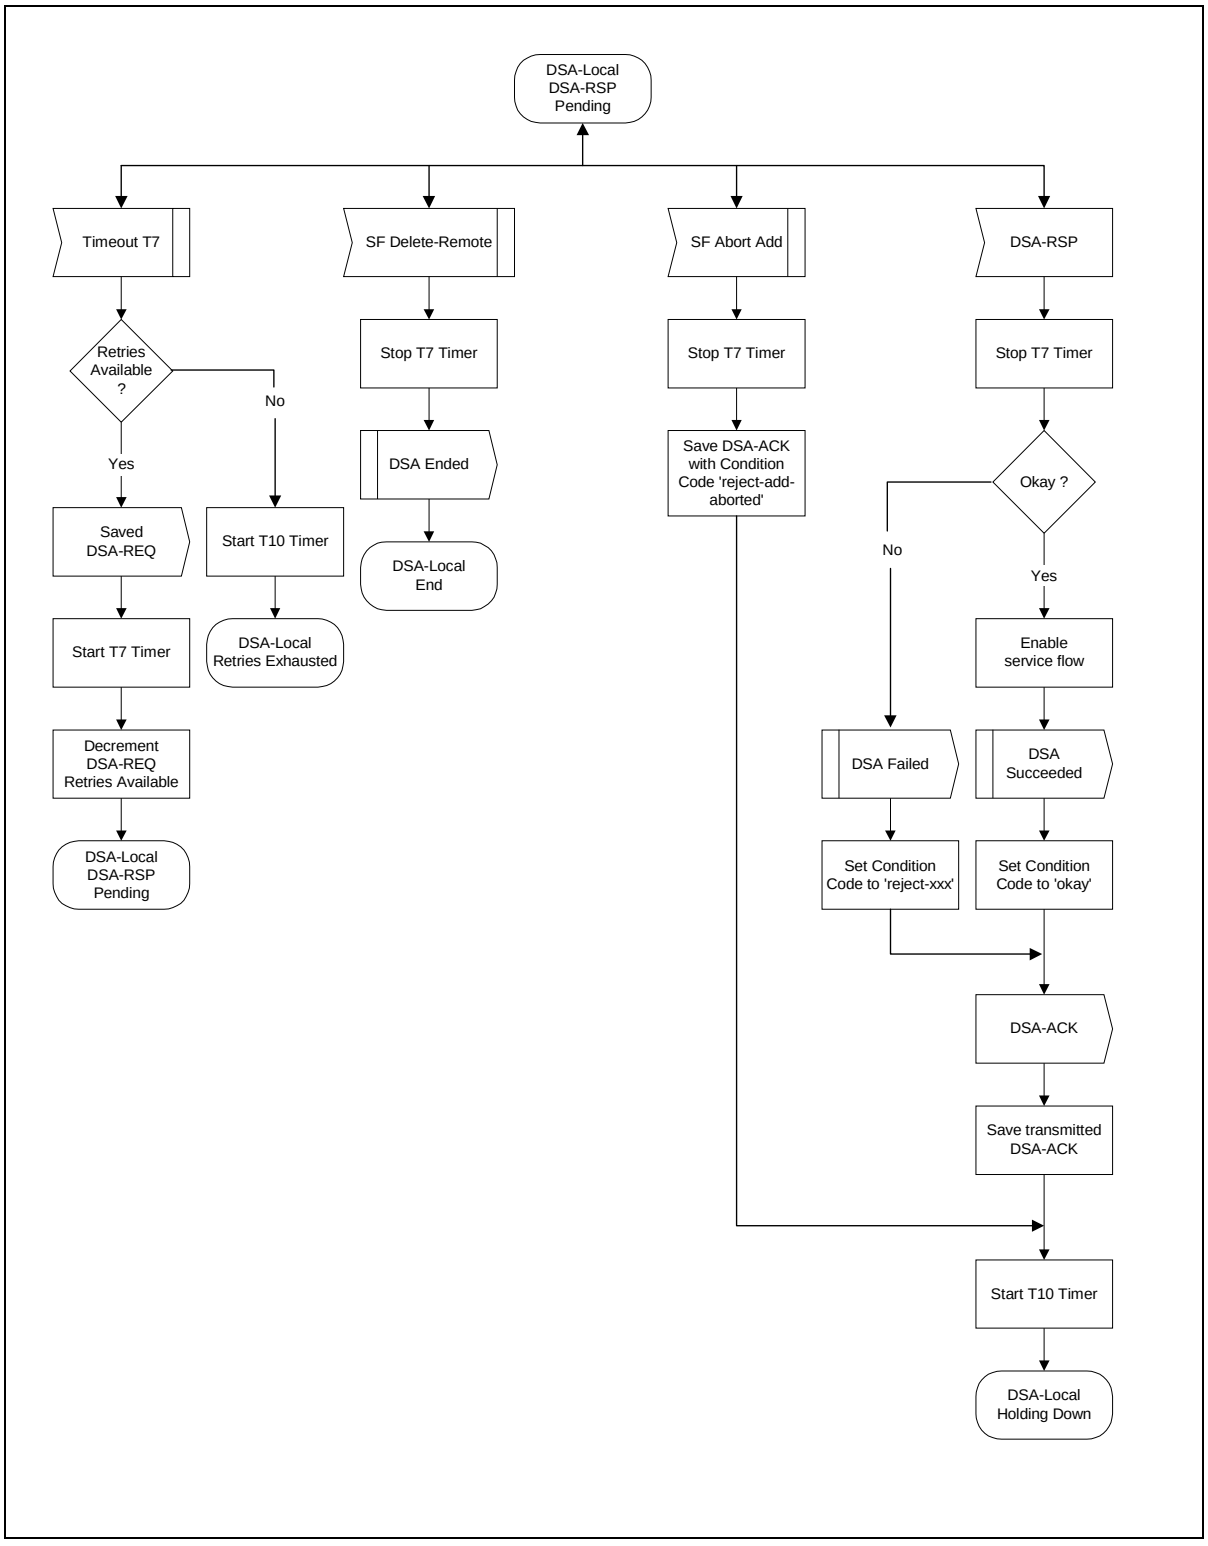


Figure 9-31. DSA—Locally Initiated Transaction DSA-RSP Pending State Flow Diagram

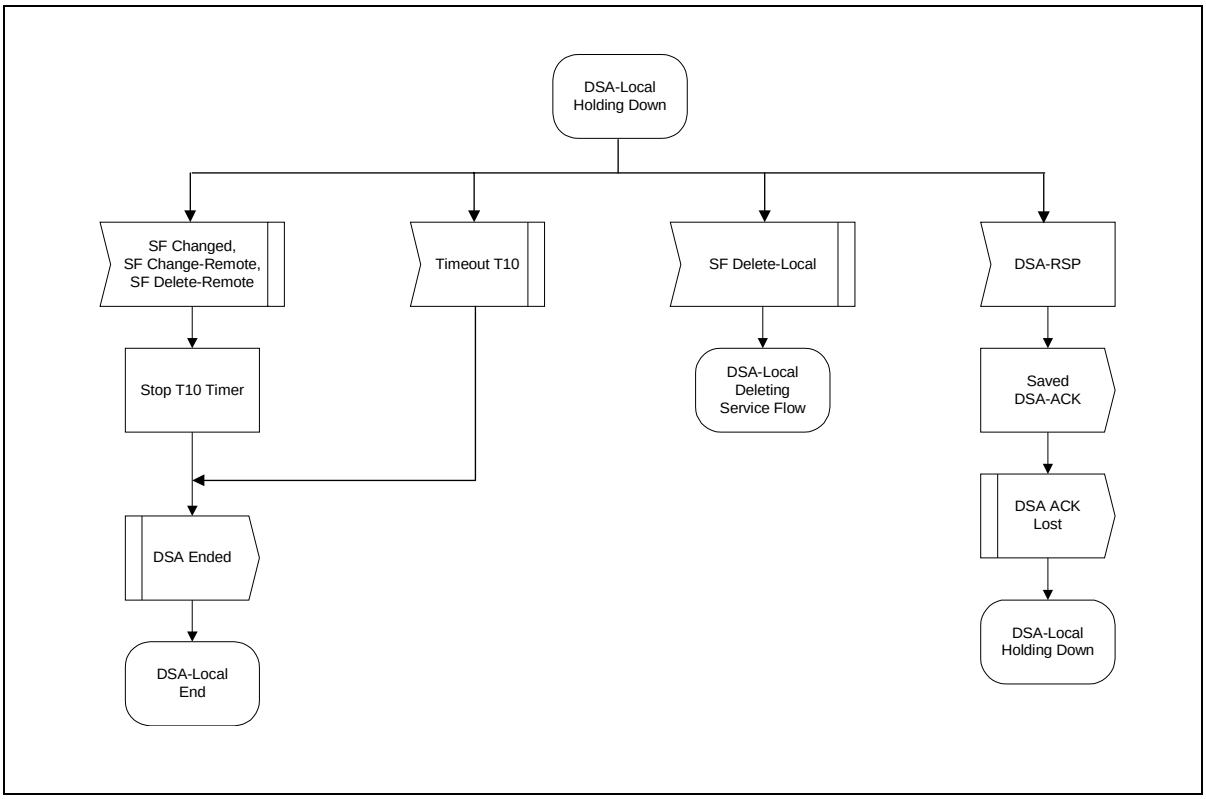


Figure 9-32. DSA—Locally Initiated Transaction Holding State Flow Diagram

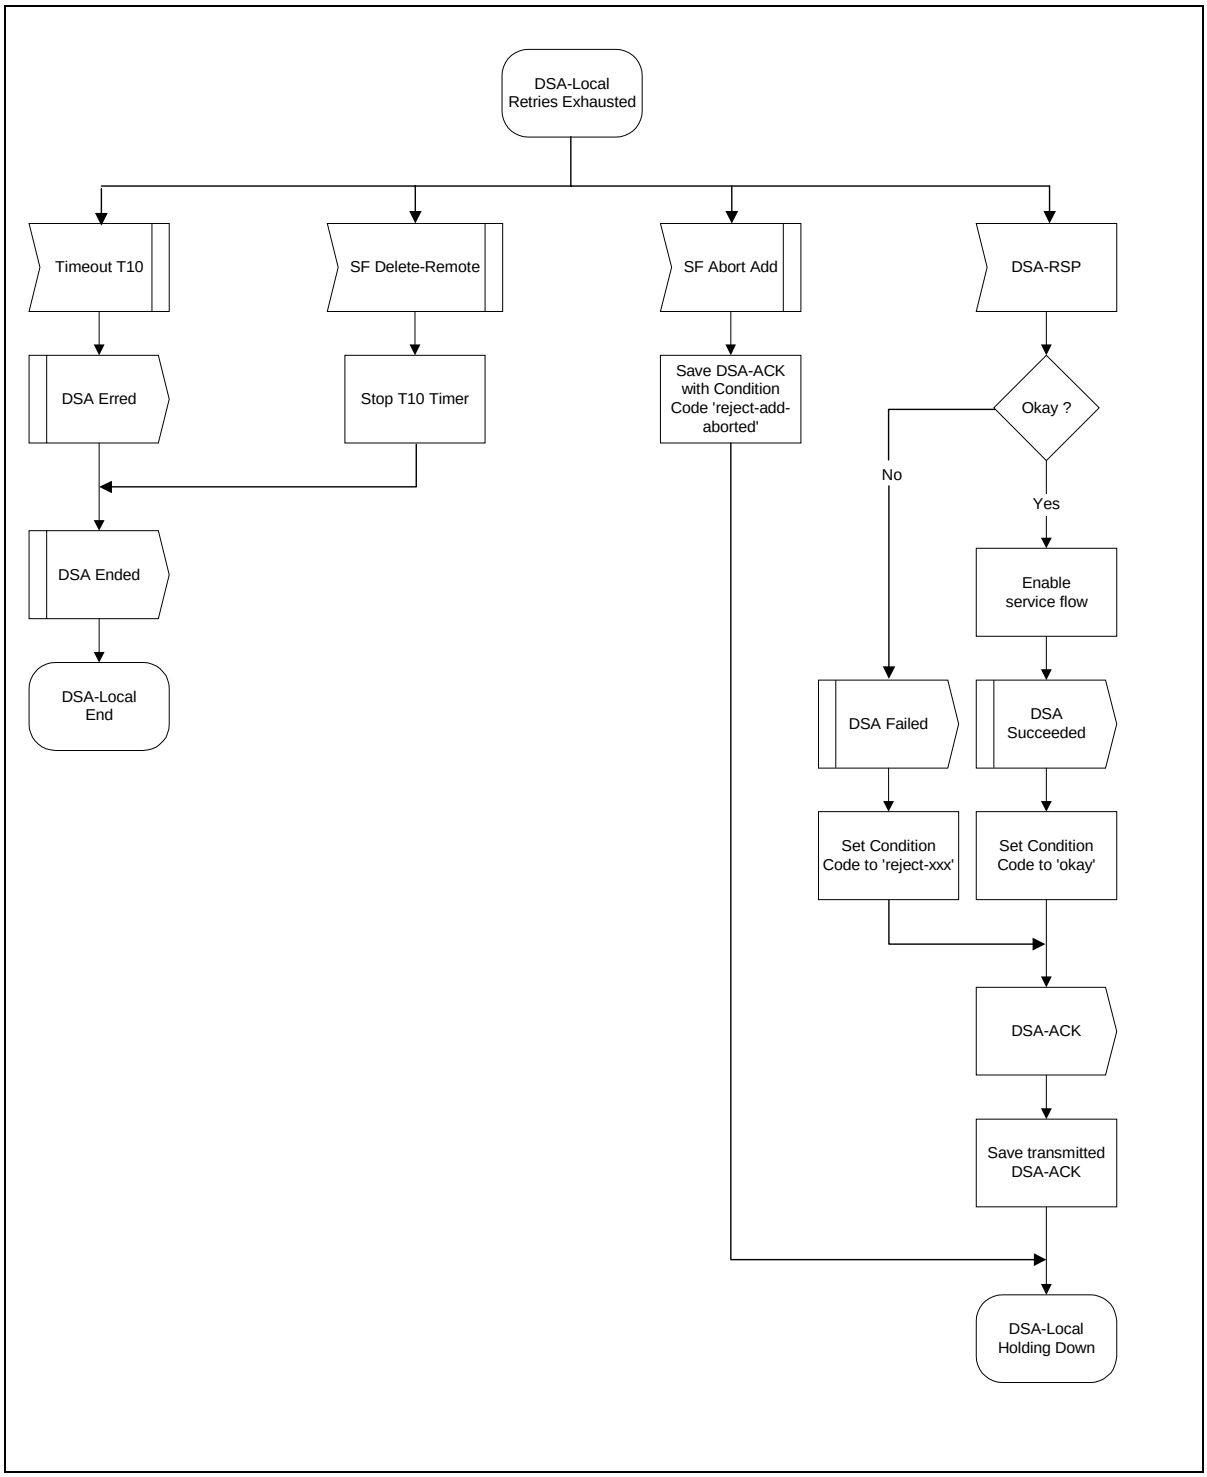


Figure 9-33. DSA—Locally Initiated Transaction Retries Exhausted State Flow Diagram

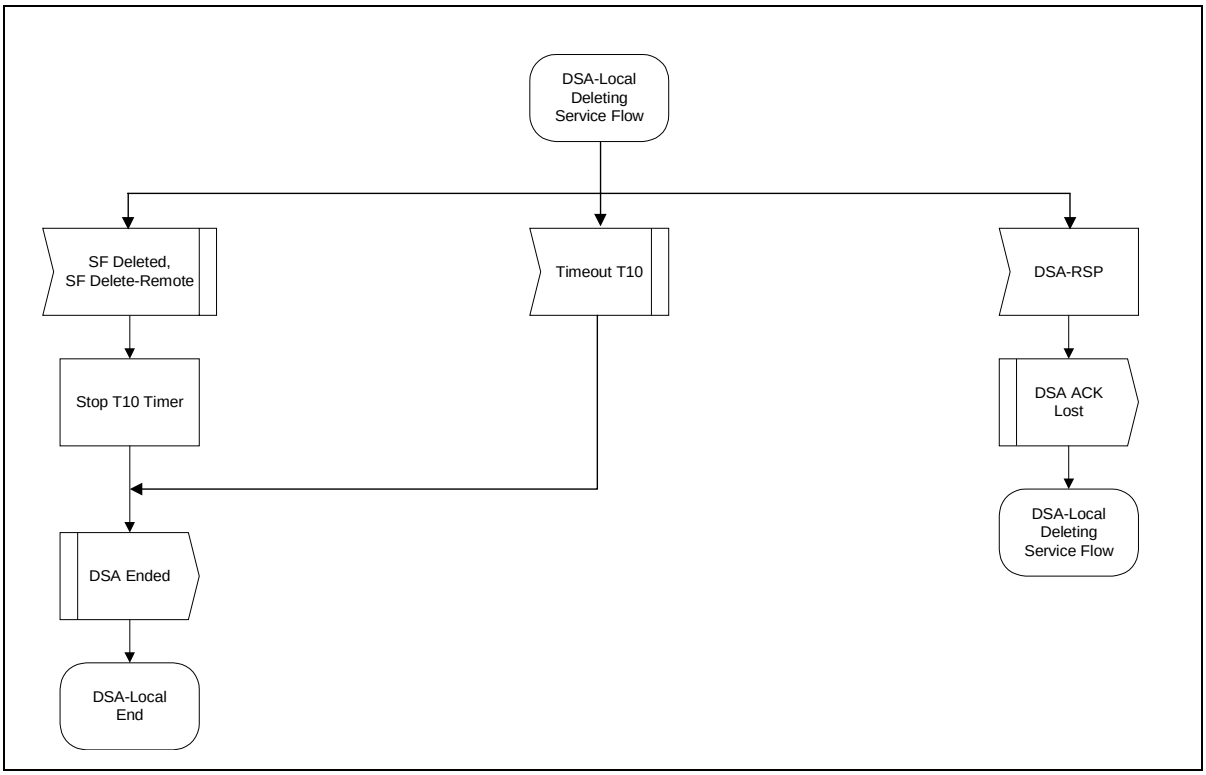


Figure 9-34. DSA—Locally Initiated Transaction Deleting Service Flow State Flow Diagram

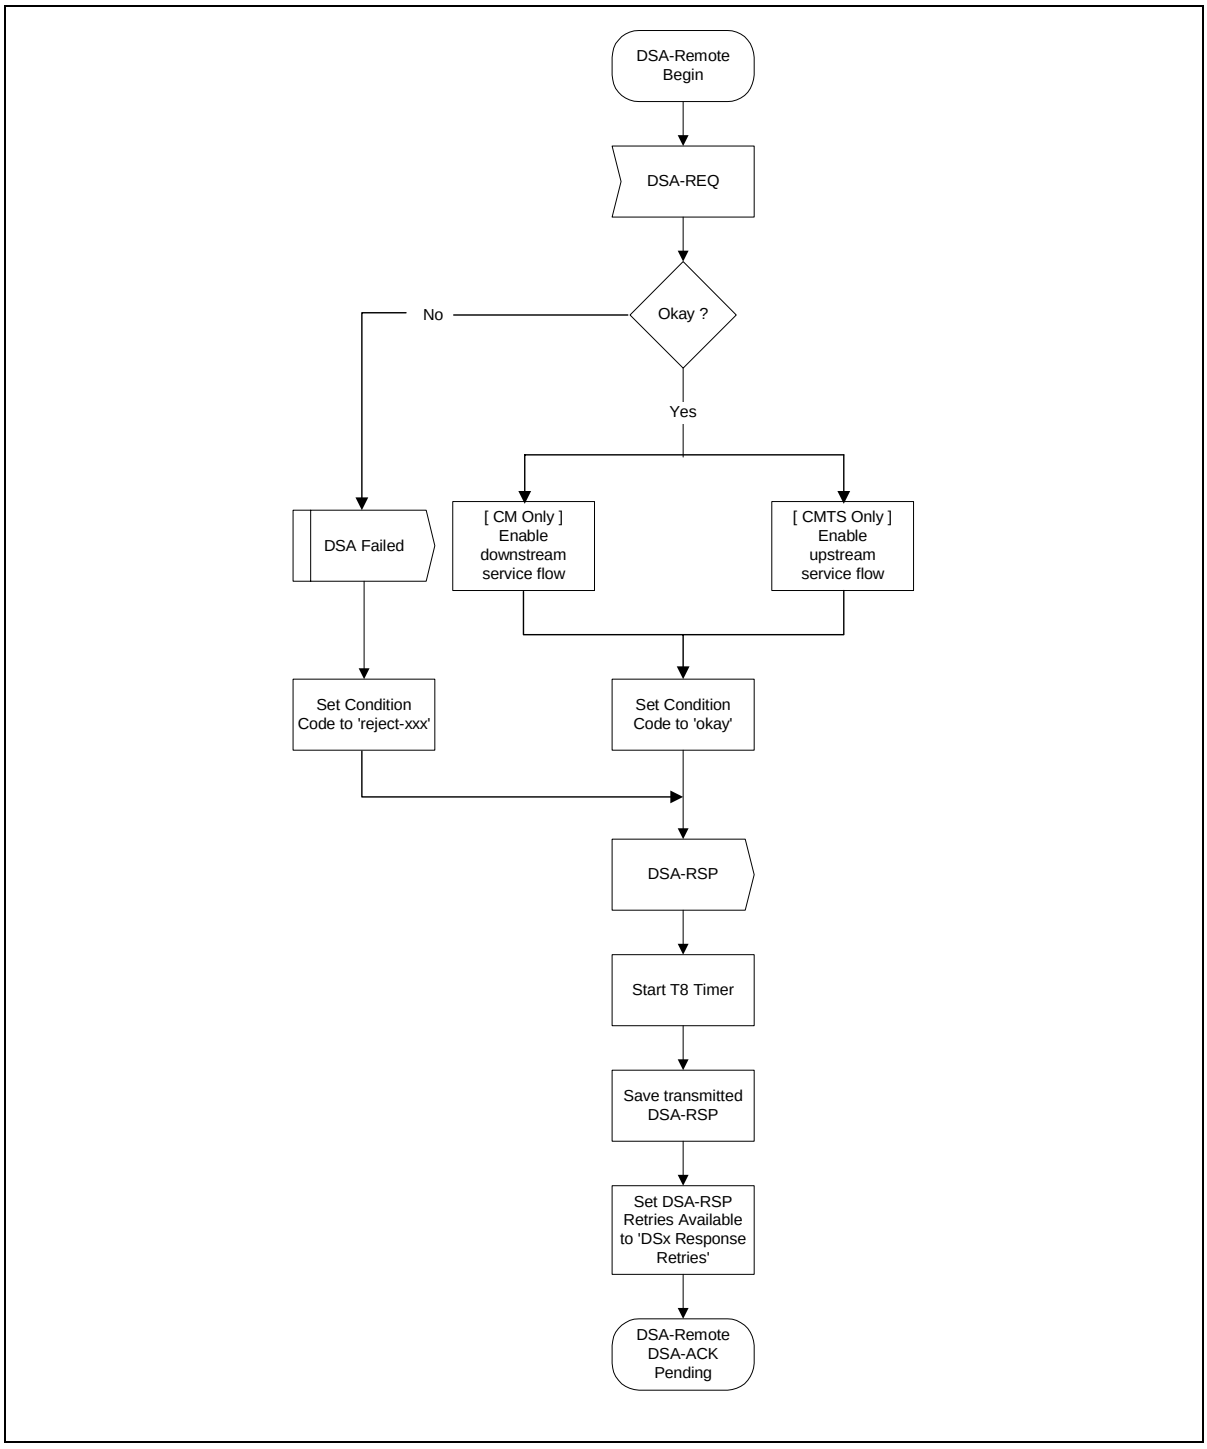


Figure 9-35. DSA—Remotely Initiated Transaction Begin State Flow Diagram

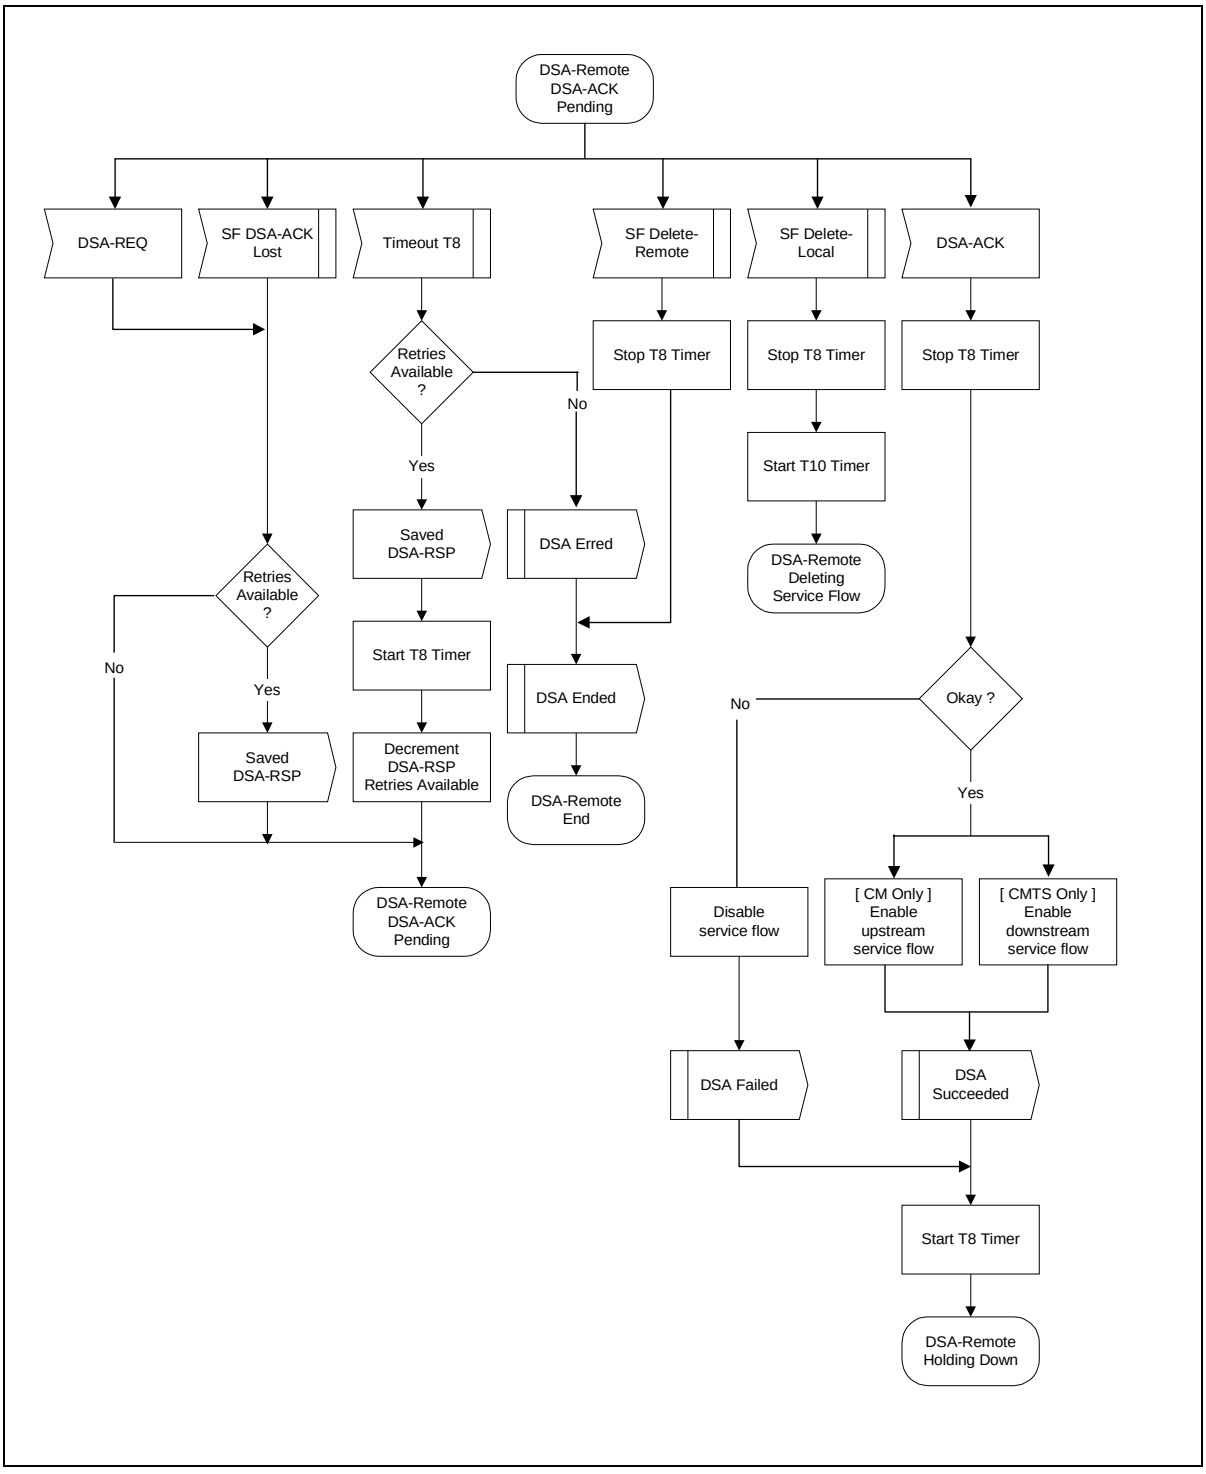


Figure 9-36. DSA—Remotely Initiated Transaction DSA-ACK Pending State Flow Diagram

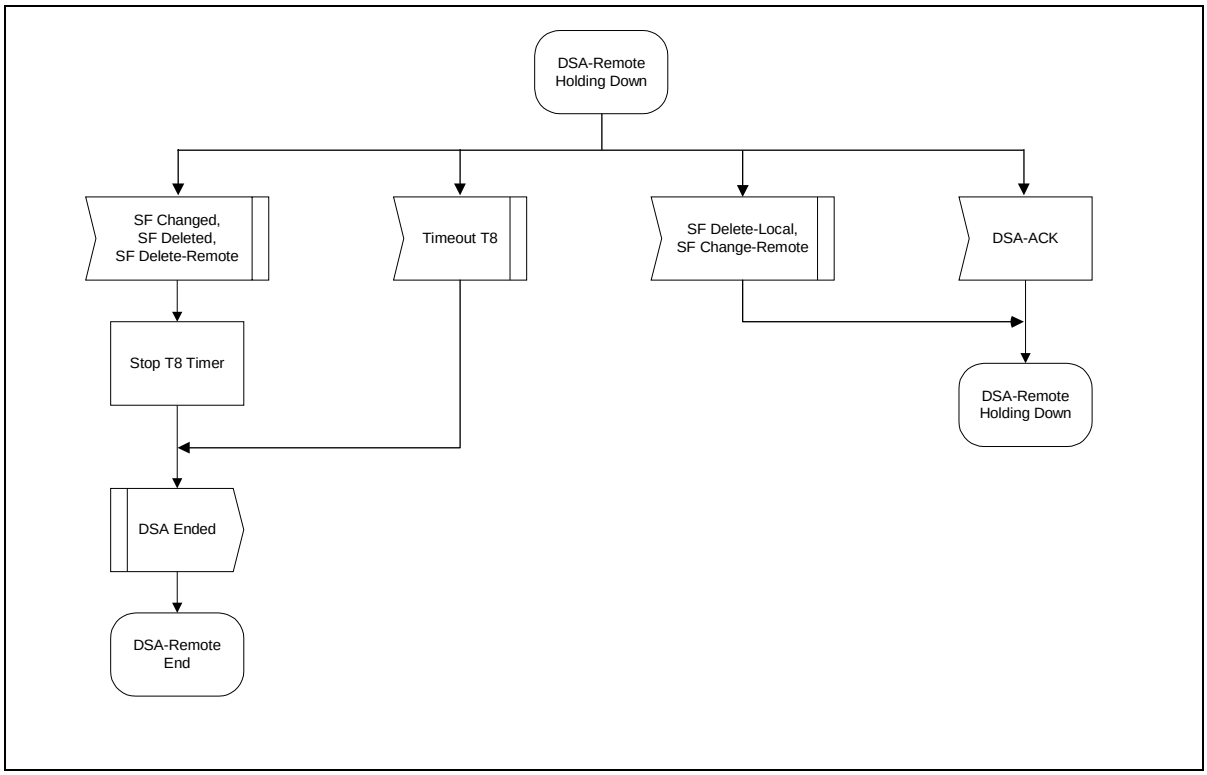


Figure 9-37. DSA—Remotely Initiated Transaction Holding Down State Flow Diagram

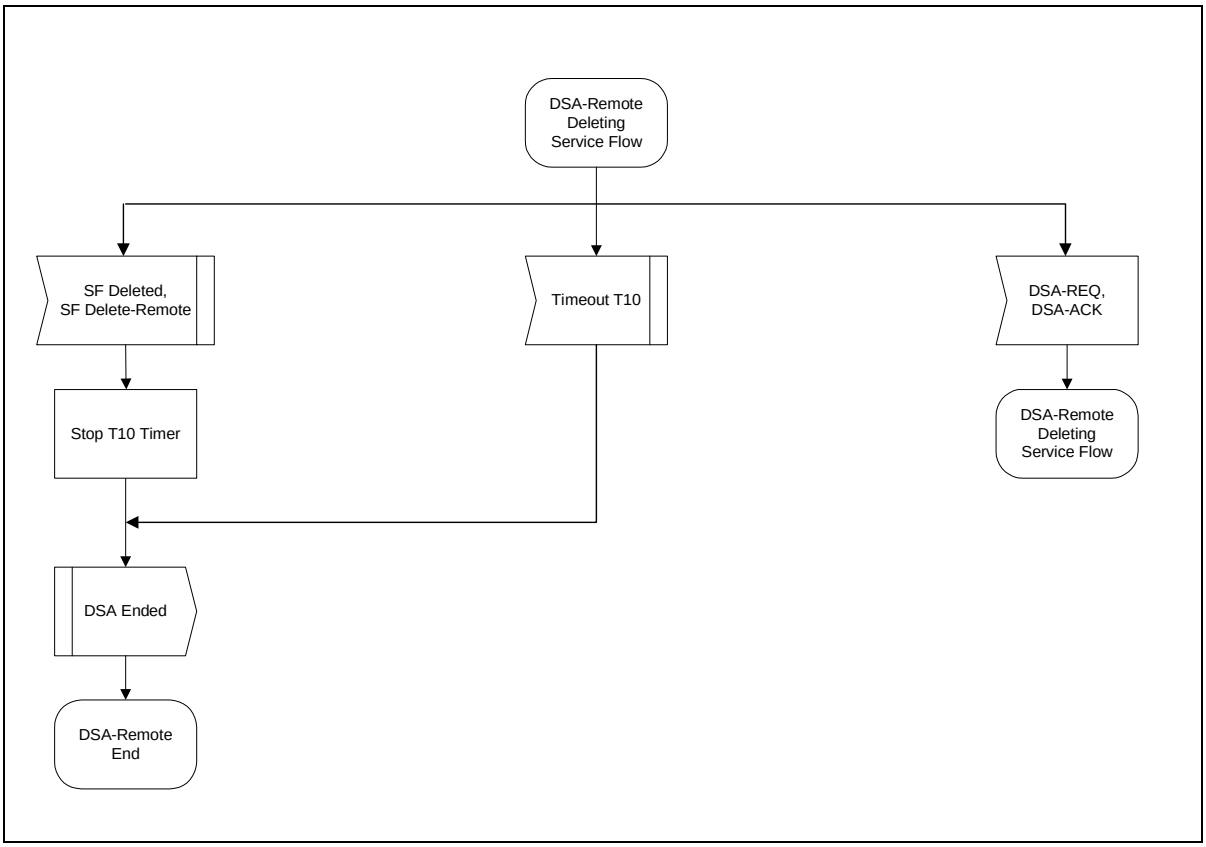


Figure 9-38. DSA—Remotely Initiated Transaction Deleting Service State Flow Diagram

9.4.3 Dynamic Service Change

The Dynamic Service Change (DSC) set of messages is used to modify the flow parameters associated with a Service Flow. Specifically, DSC can:

- Modify the Service Flow Specification
- Add, Delete or Replace a Flow Classifier
- Add, Delete or Set PHS elements

A single DSC message exchange can modify the parameters of one downstream service flow and/or one upstream service flow.

To prevent packet loss, any required bandwidth change must be sequenced between the application generating the data and the bandwidth parameters of the Service Flow carrying the data. Because MAC messages can be lost, the timing of Service Flow parameter changes can vary, and it occurs at different times in the CM and CMTS. Applications should reduce their transmitted data bandwidth before initiating a DSC to reduce the Service Flow bandwidth, and should not increase their transmitted data bandwidth until after the completion of a DSC increasing the Service Flow bandwidth.

The CMTS controls both upstream and downstream scheduling. Scheduling is based on data transmission requests and is subject to the limits contained in the current Service Flow parameters at the CMTS. The timing of Service Flow parameter changes, and any consequent scheduling changes, is independent of both direction and

whether there is an increase or decrease in bandwidth. The CMTS always changes Service Flow parameters on receipt of a DSC-REQ (CM-initiated transaction) or DSC-RSP (CMTS-initiated transaction).

The CMTS also controls the downstream transmit behavior. The change in downstream transmit behavior is always coincident with the change in downstream scheduling (*i.e.*, CMTS controls both and changes both simultaneously).

The CM controls the upstream transmit requests, subject to limits contained in the current Service Flow parameters at the CM. The timing of Service Flow parameter changes in the CM, and any consequent CM transmit request behavior changes, is a function of which device initiated the transaction. For a CM-initiated DSC-REQ, the Service Flow parameters are changed on receipt of the DSC-RSP from the CMTS. For a CMTS-initiated DSC-REQ, the Service Flow parameters are changed on receipt of the DSC-REQ from the CMTS.

Any service flow can be deactivated with a Dynamic Service Change command by sending a DSC-REQ message, referencing the Service Flow Identifier, and including a null ActiveQoSParameterSet. However, if a Primary Service Flow of a CM is deactivated that CM is de-registered and MUST re-register. Therefore, care should be taken before deactivating such Service Flows. If a Service Flow that was provisioned during registration is deactivated, the provisioning information for that Service Flow MUST be maintained until the Service Flow is reactivated.

A CM MUST have only one DSC transaction outstanding per Service Flow. If it detects a second transaction initiated by the CMTS, the CM MUST abort the transaction it initiated and allow the CMTS initiated transaction to complete.

A CMTS MUST have only one DSC transaction outstanding per Service Flow. If it detects a second transaction initiated by the CM, the CMTS MUST abort the transaction the CM initiated and allow the CMTS initiated transaction to complete.

Note: Currently anticipated applications would probably control a Service Flow through either the CM or CMTS, and not both. Therefore the case of a DSC being initiated simultaneously by the CM and CMTS is considered as an exception condition and treated as one.

9.4.3.1 CM-Initiated Dynamic Service Change

A CM that needs to change a Service Flow definition performs the following operations.

The CM informs the CMTS using a Dynamic Service Change Request message (DSC-REQ). The CMTS MUST decide if the referenced Service Flow can support this modification. The CMTS MUST respond with a Dynamic Service Change Response (DSC-RSP) indicating acceptance or rejection. The CM reconfigures the Service Flow if appropriate, and then MUST respond with a Dynamic Service Change Acknowledge (DSC-ACK).

CMTS		**CM**
		Service Flow Requires Modifying
Receive DSC-REQ	**<--------- DSC-REQ ----------**	**Send DSC-REQ**
Validate Request		
Modify Service Flow		
Send DSC-RSP	**---------- DSC-RSP --------->**	**Receive DSC-RSP**
		Modify Service Flow
Receive DSC-ACK	**<--------- DSC-ACK ----------**	**Send DSC-ACK**

Figure 9-39. CM-Initiated DSC

9.4.3.2 CMTS-Initiated Dynamic Service Change

A CMTS that needs to change a Service Flow definition performs the following operations.

The CMTS MUST decide if the referenced Service Flow can support this modification. If so, the CMTS informs the CM using a Dynamic Service Change Request message (DSC-REQ). The CM checks that it can support the service change, and MUST respond using a Dynamic Service Change Response (DSC-RSP) indicating acceptance or rejection. The CMTS reconfigures the Service Flow if appropriate, and then MUST respond with a Dynamic Service Change Acknowledgment (DSC-ACK)

CMTS		CM
Service Flow Requires Modifying		
Send DSC-REQ	**---------- DSC-REQ --------->**	**Receive DSC-REQ**
		Modify Service Flow
Receive DSC-RSP	**<--------- DSC-RSP ----------**	**Send DSC-RSP**
Modify Service Flow		
Send DSC-ACK	**---------- DSC-ACK --------->**	**Receive DSC-ACK**

Figure 9-40. CMTS-Initiated DSC

9.4.3.3 Dynamic Service Change State Transition Diagrams

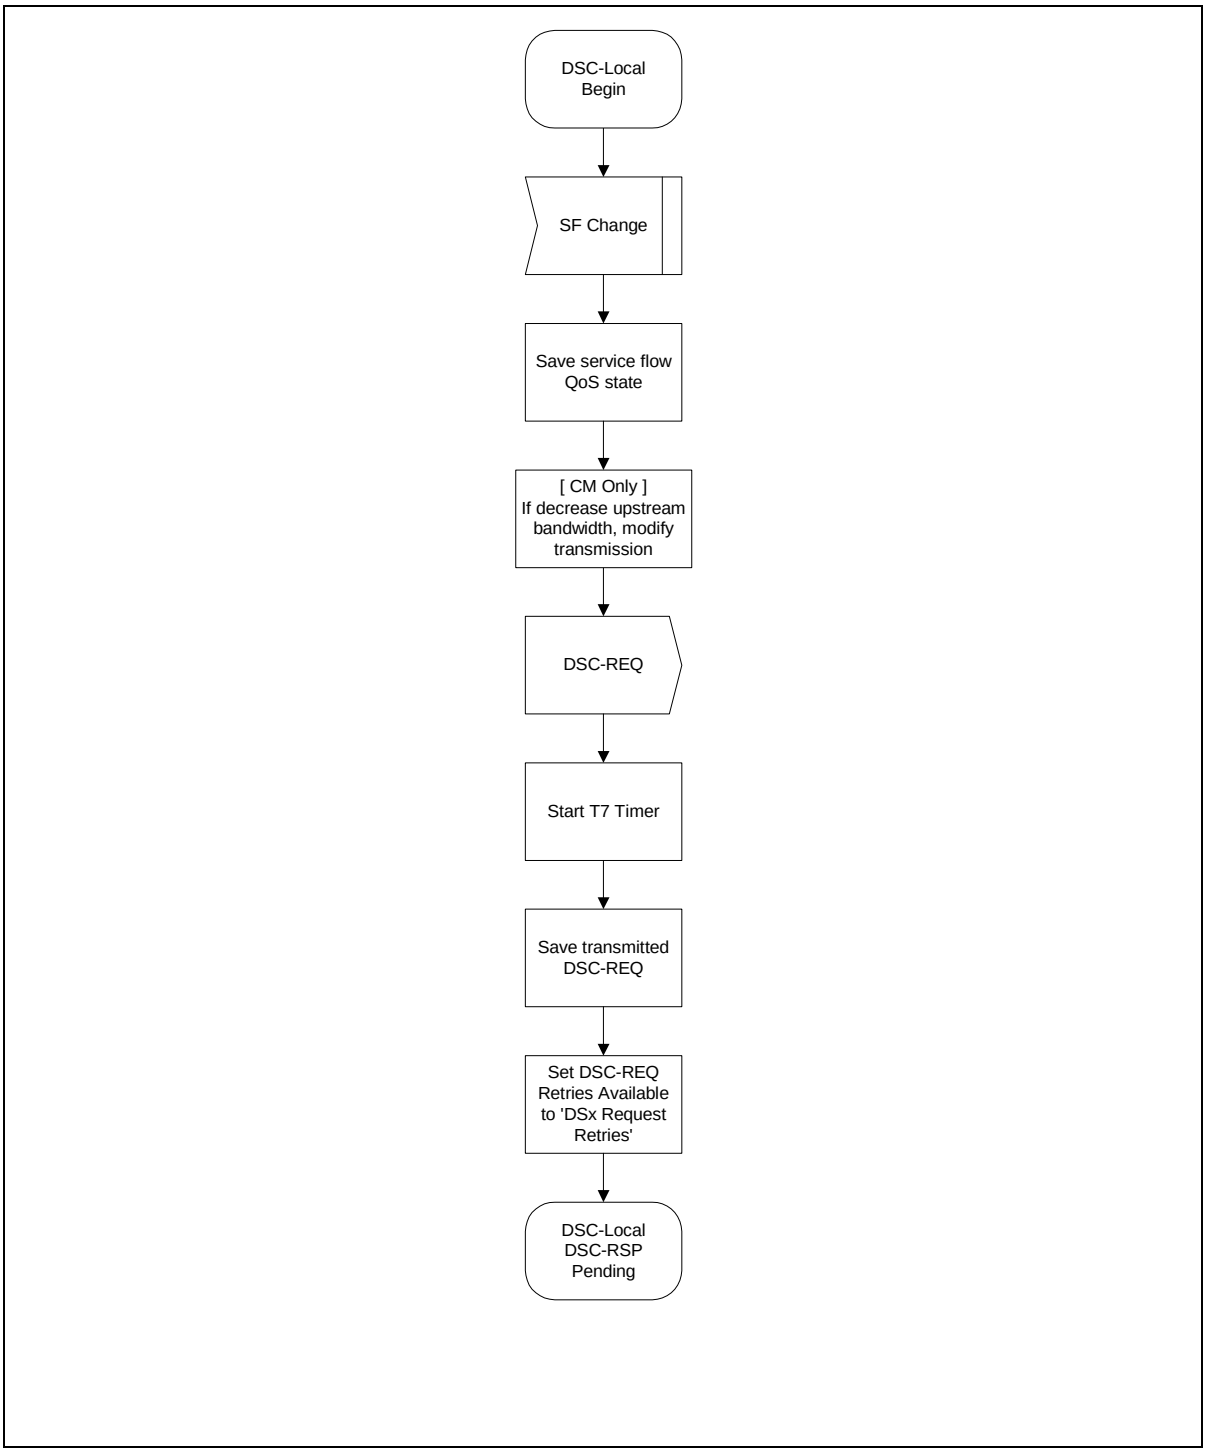


Figure 9-41. DSC—Locally Initiated Transaction Begin State Flow Diagram

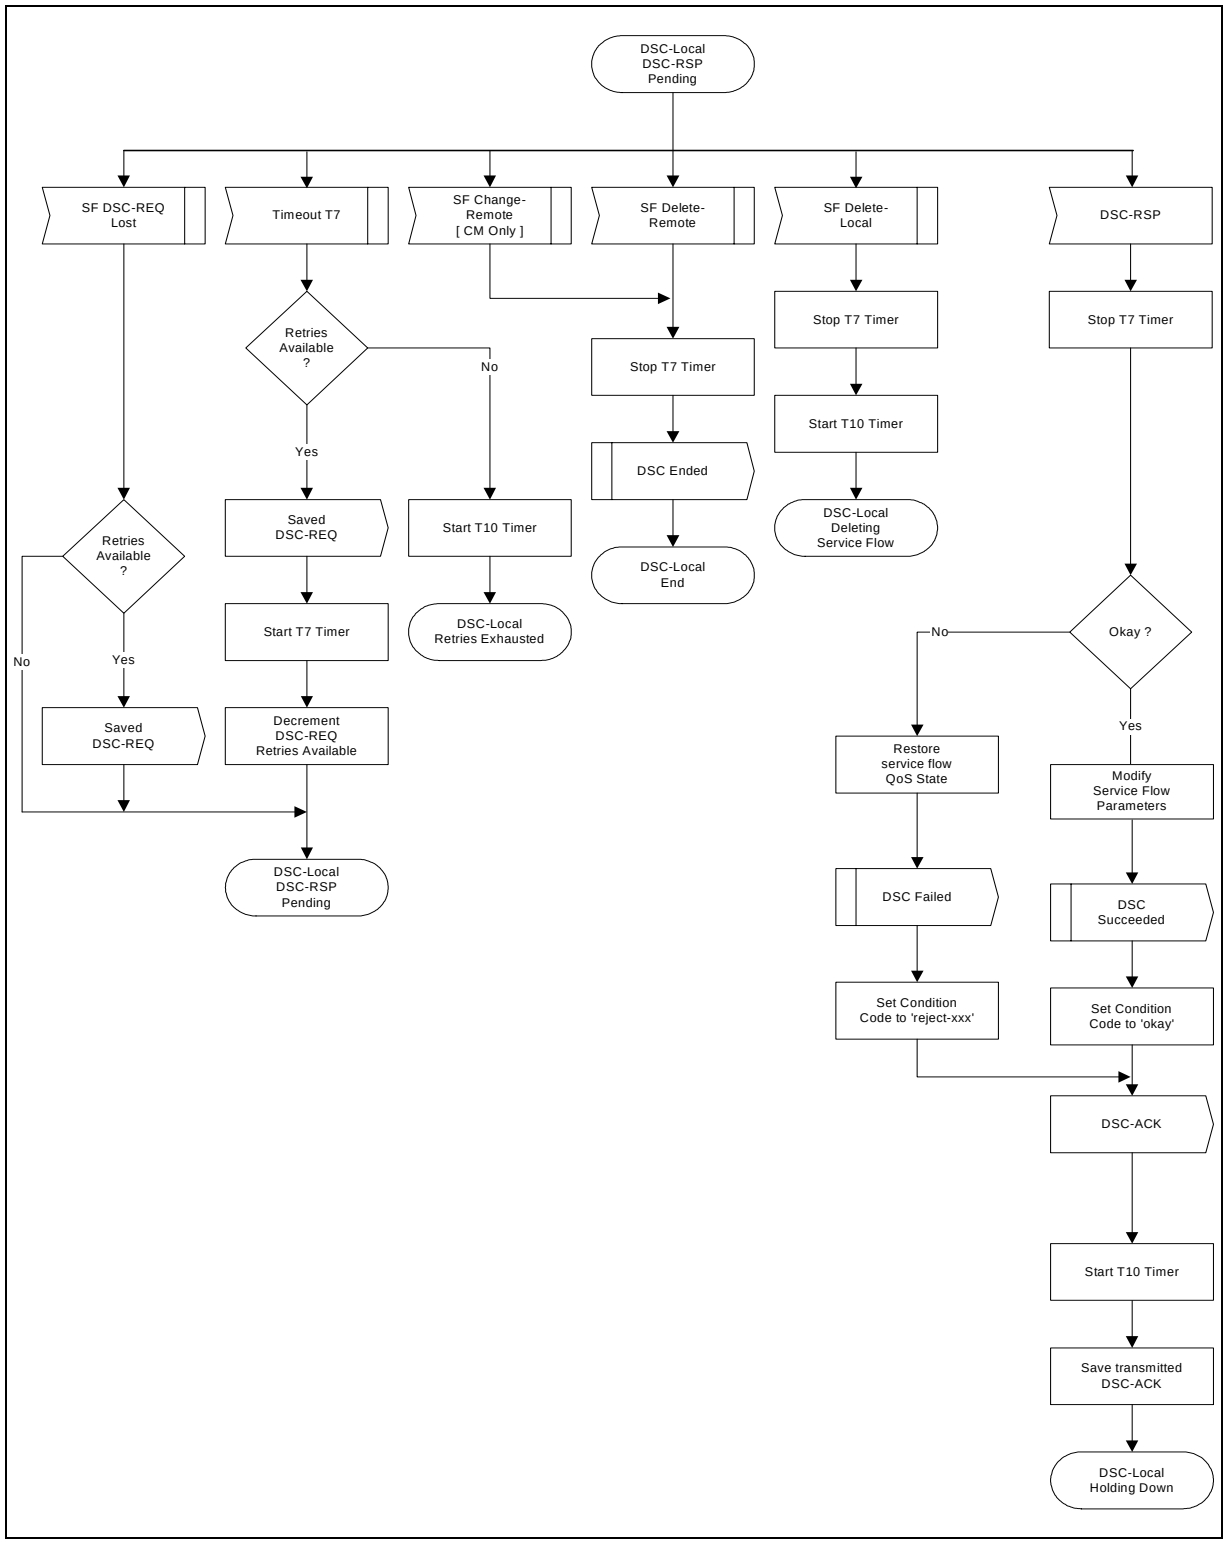


Figure 9-42. DSC—Locally Initiated Transaction DSC-RSP Pending State Flow Diagram

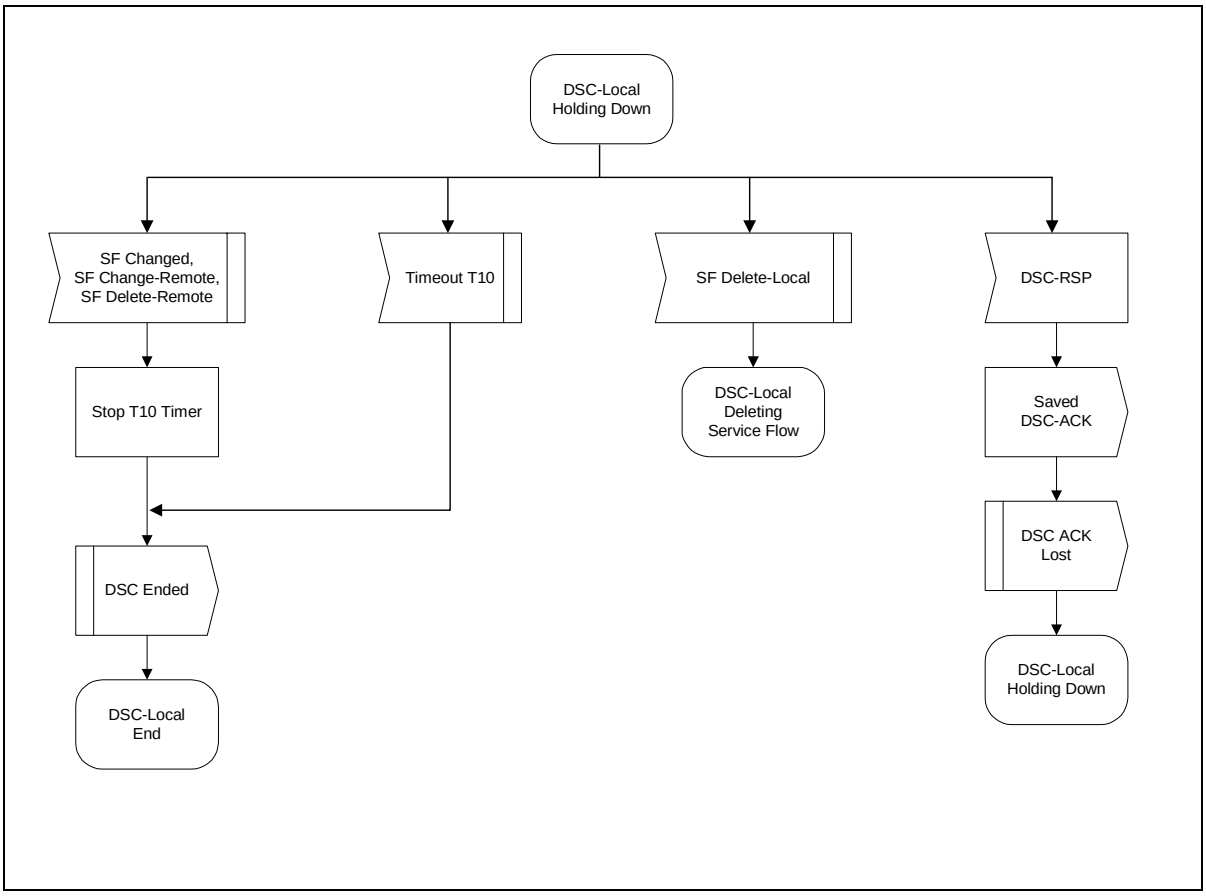


Figure 9-43. DSC—Locally Initiated Transaction Holding Down State Flow Diagram

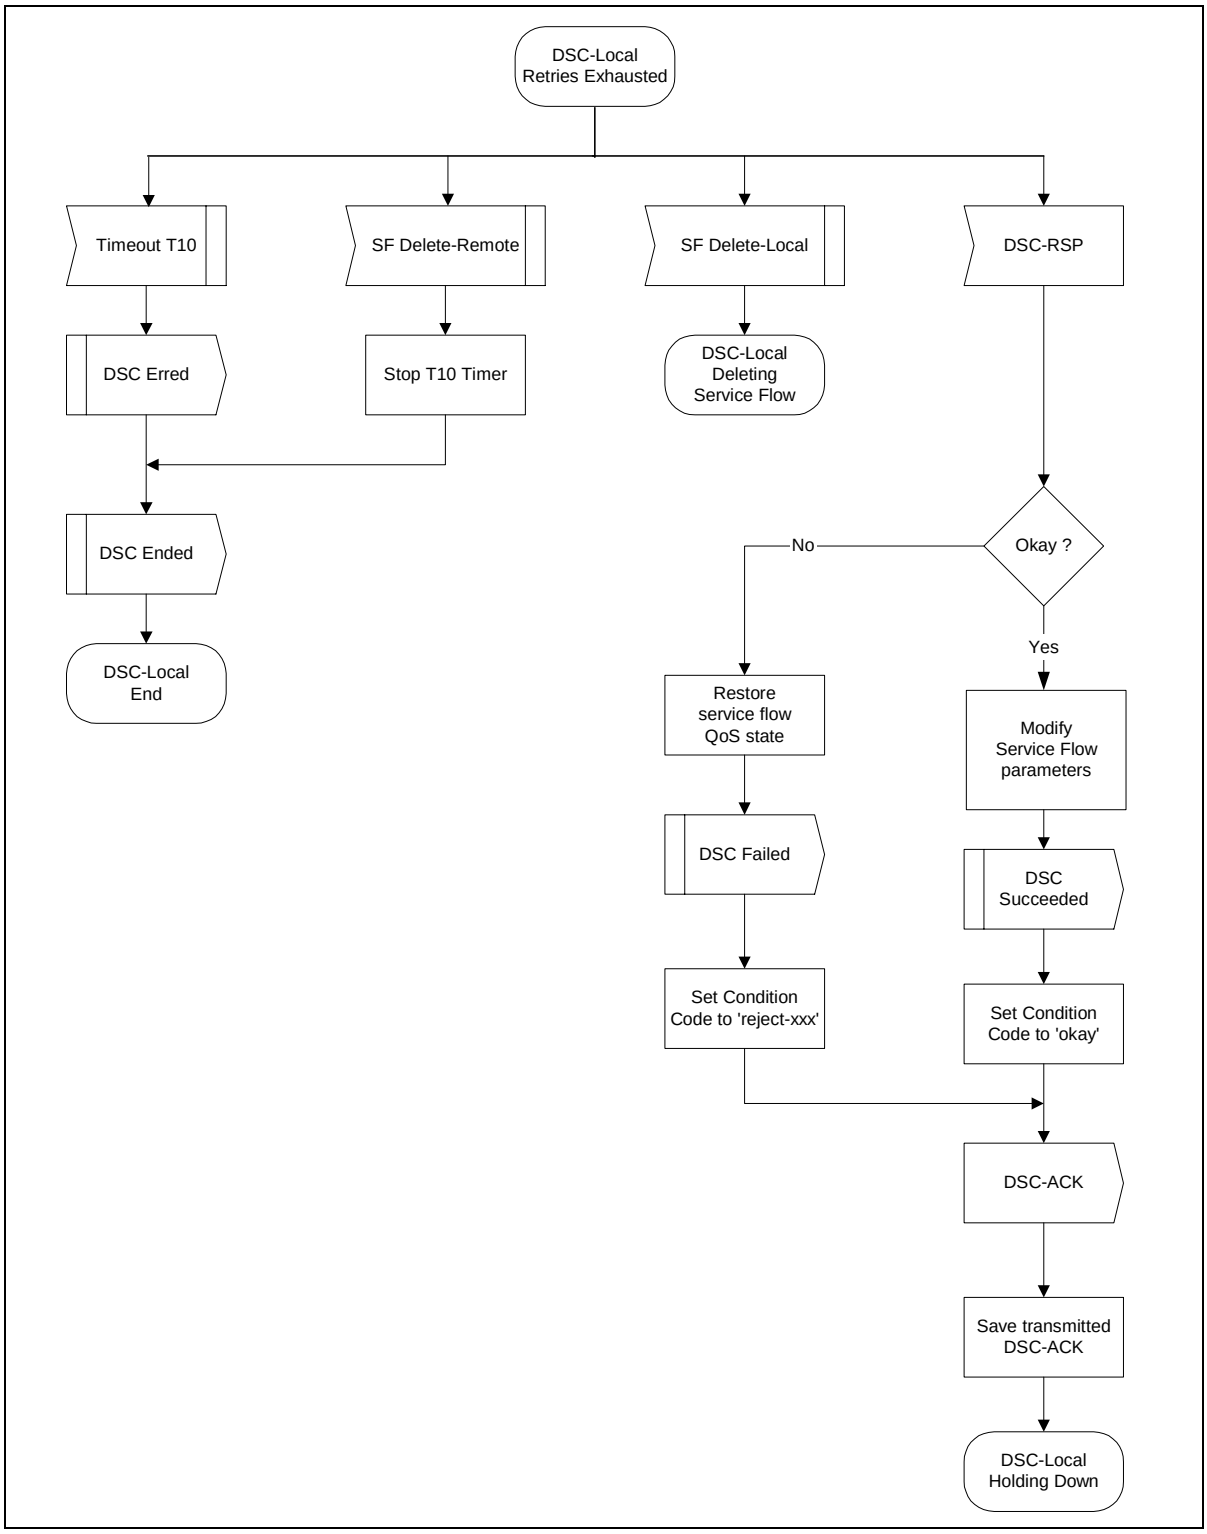


Figure 9-44. DSC—Locally Initiated Transaction Retries Exhausted State Flow Diagram

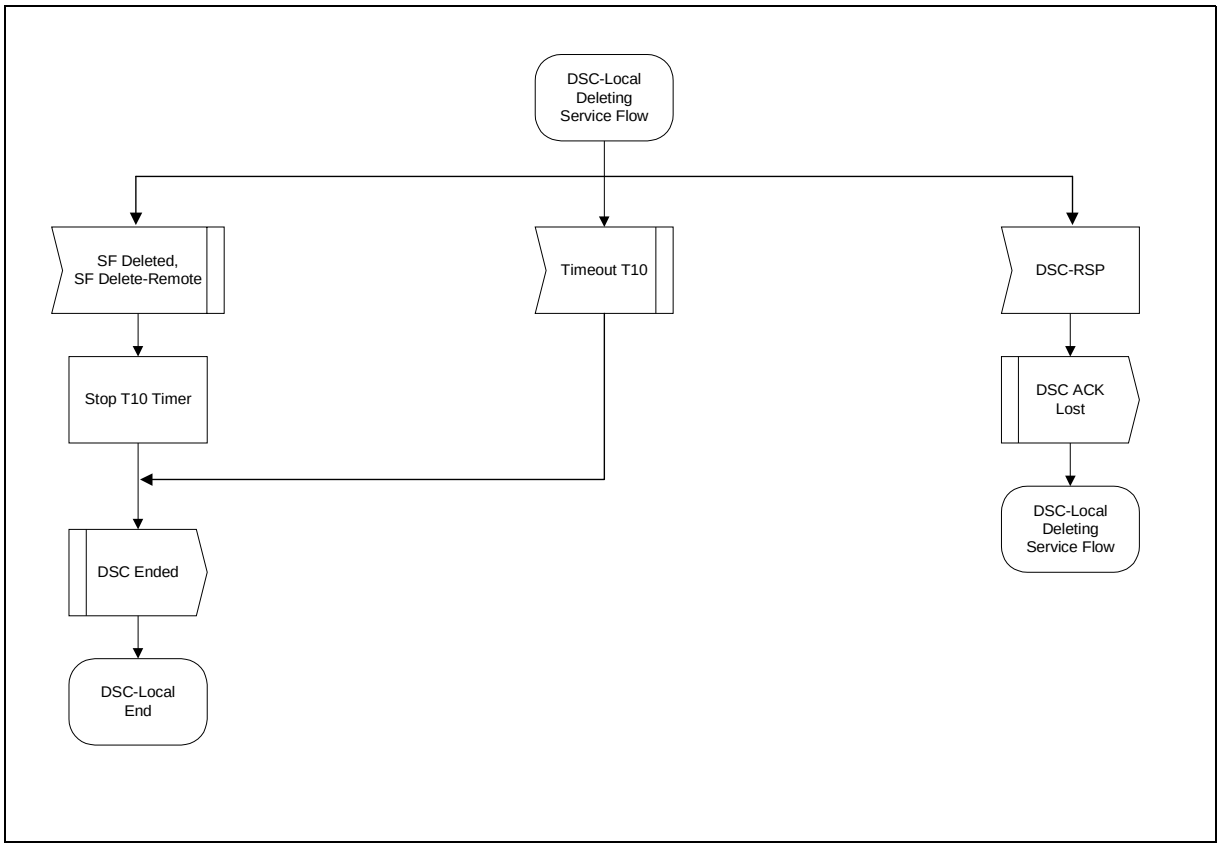


Figure 9-45. DSC—Locally Initiated Transaction Deleting Service Flow State Flow Diagram

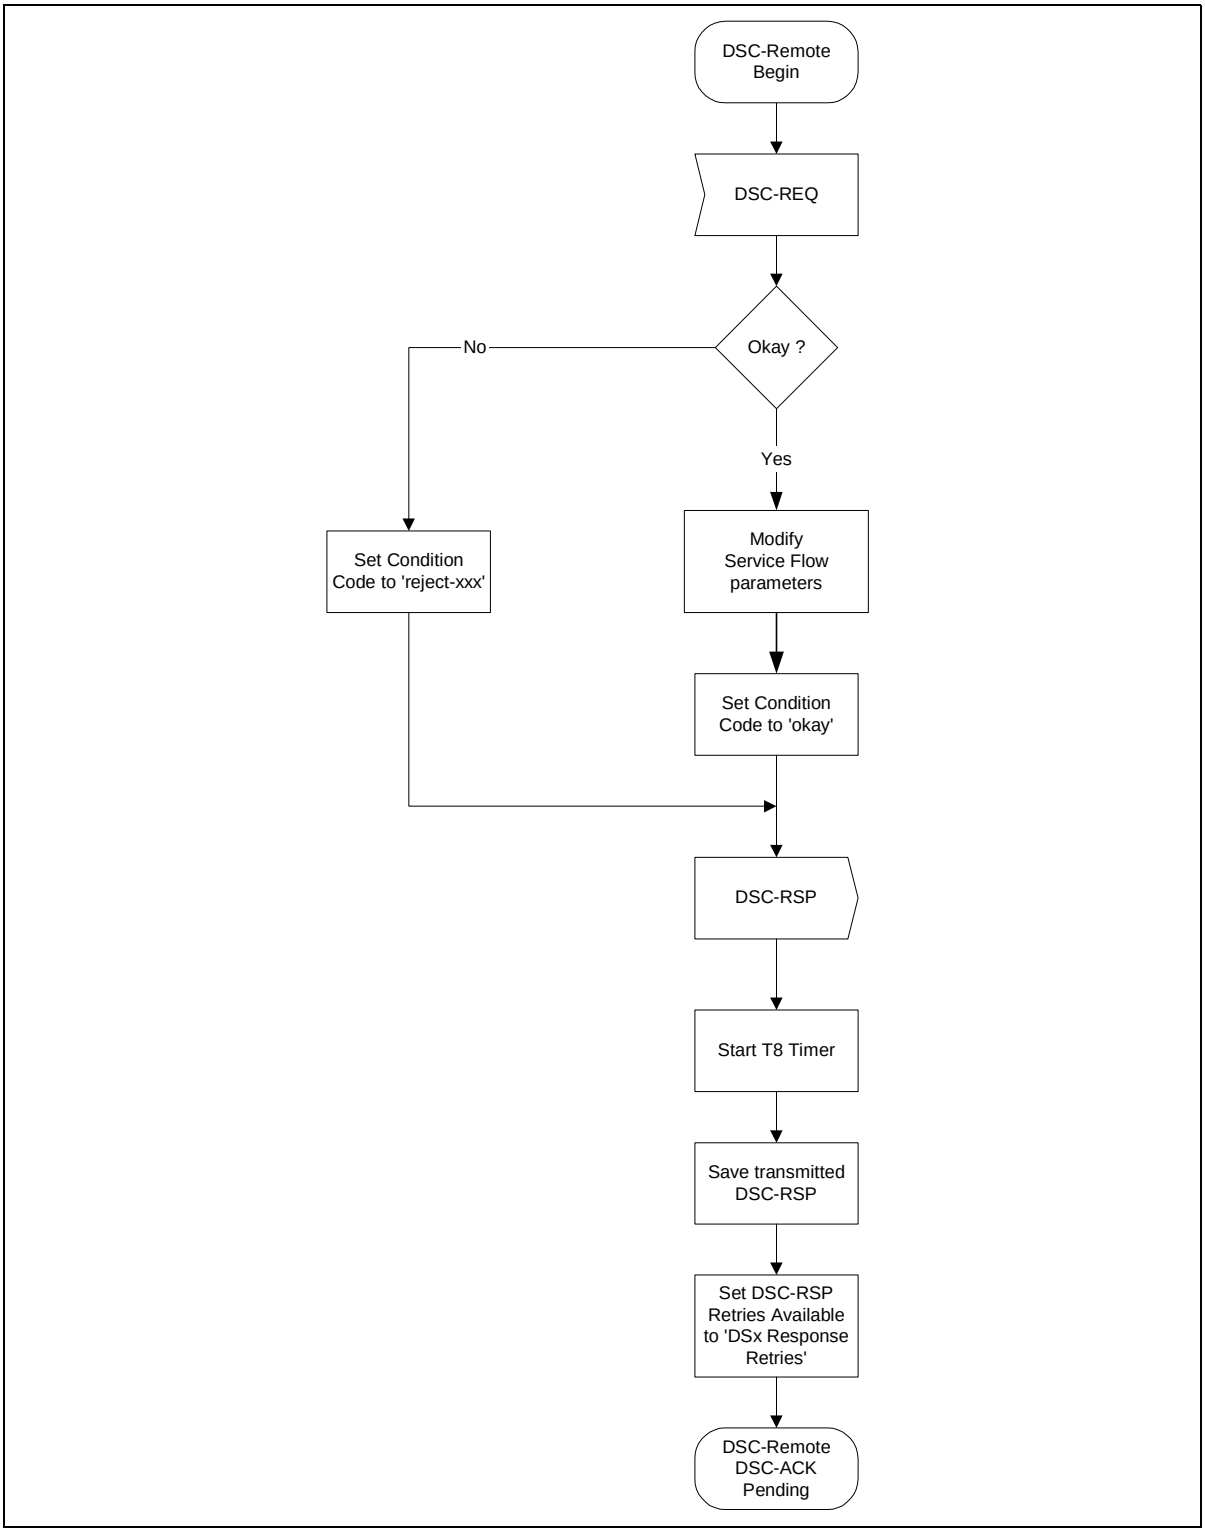


Figure 9-46. DSC—Remotely Initiated Transaction Begin State Flow Diagram

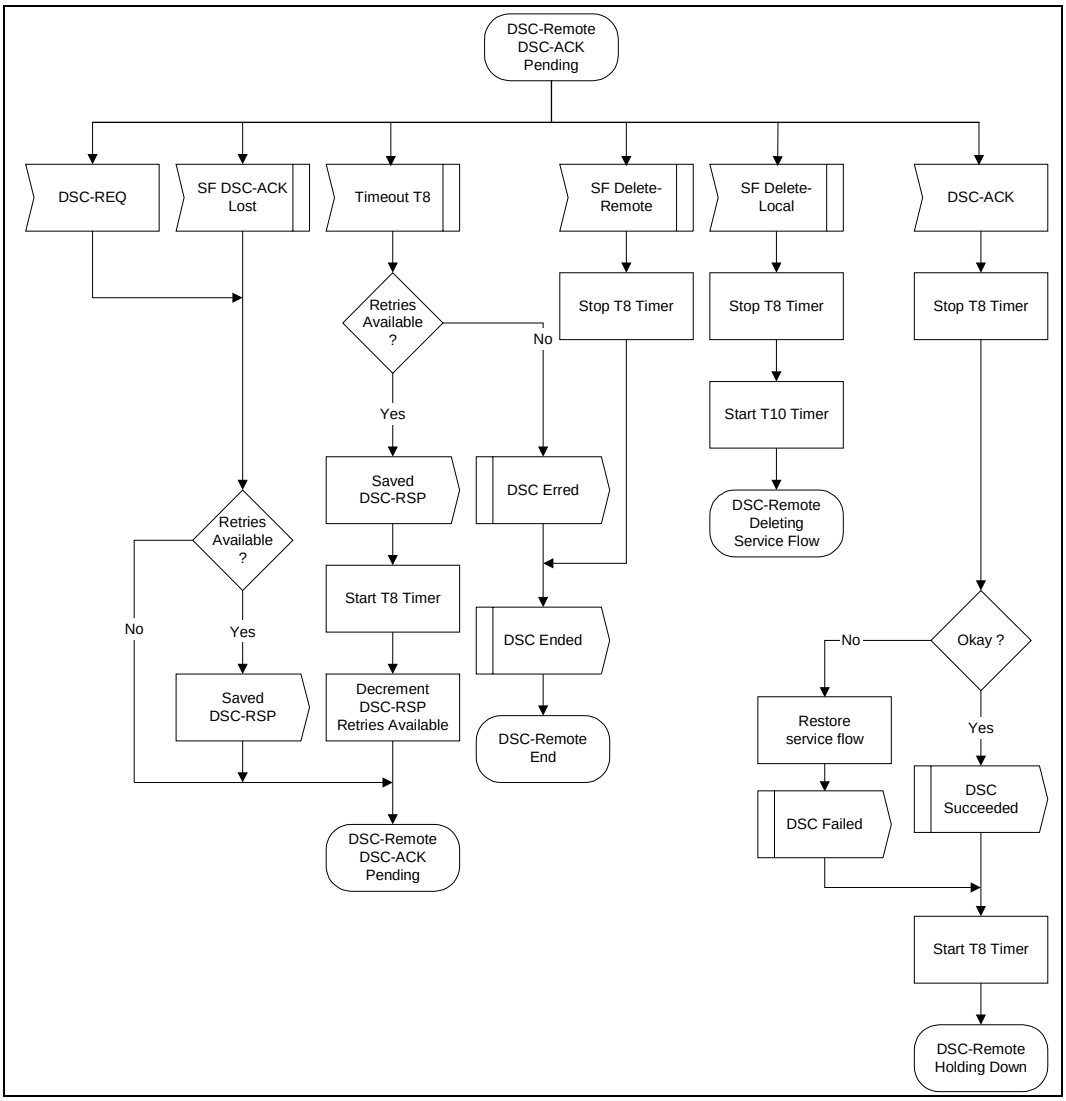


Figure 9-47. DSC—Remotely Initiated Transaction DSC-ACK Pending State Flow Diagram

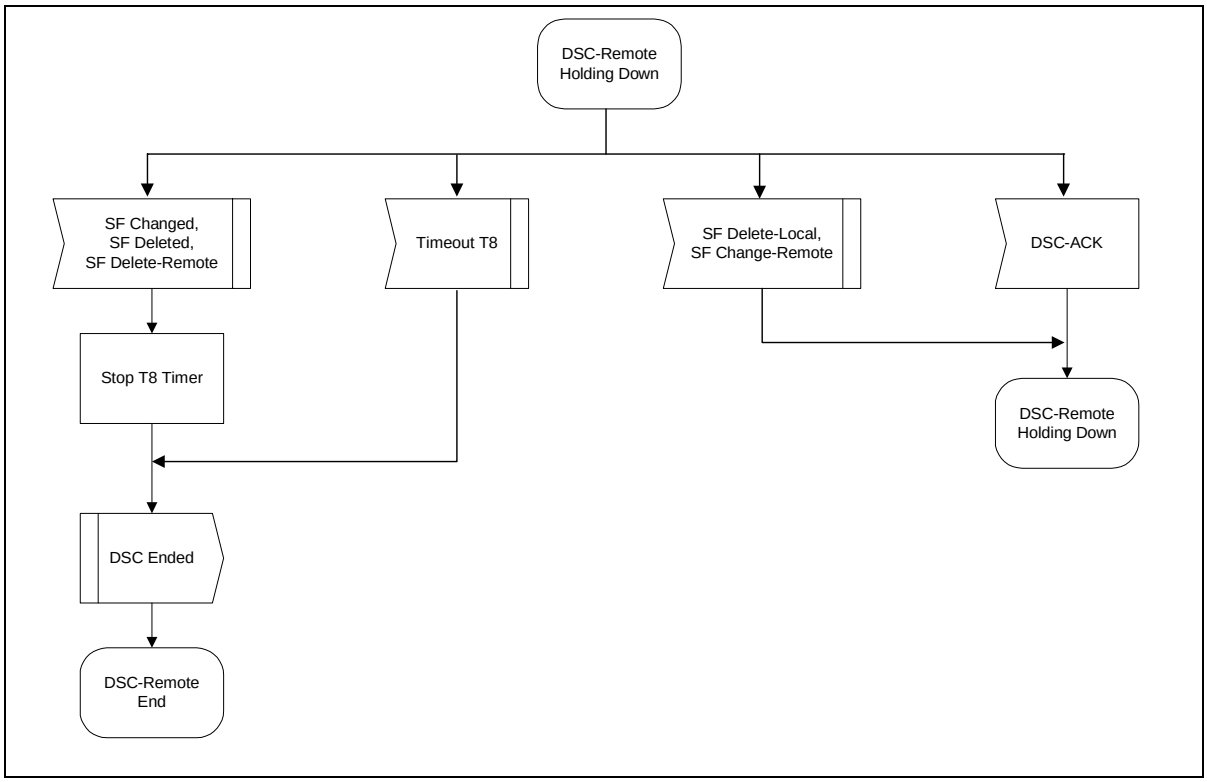


Figure 9-48. DSC—Remotely Initiated Transaction Holding Down State Flow Diagram

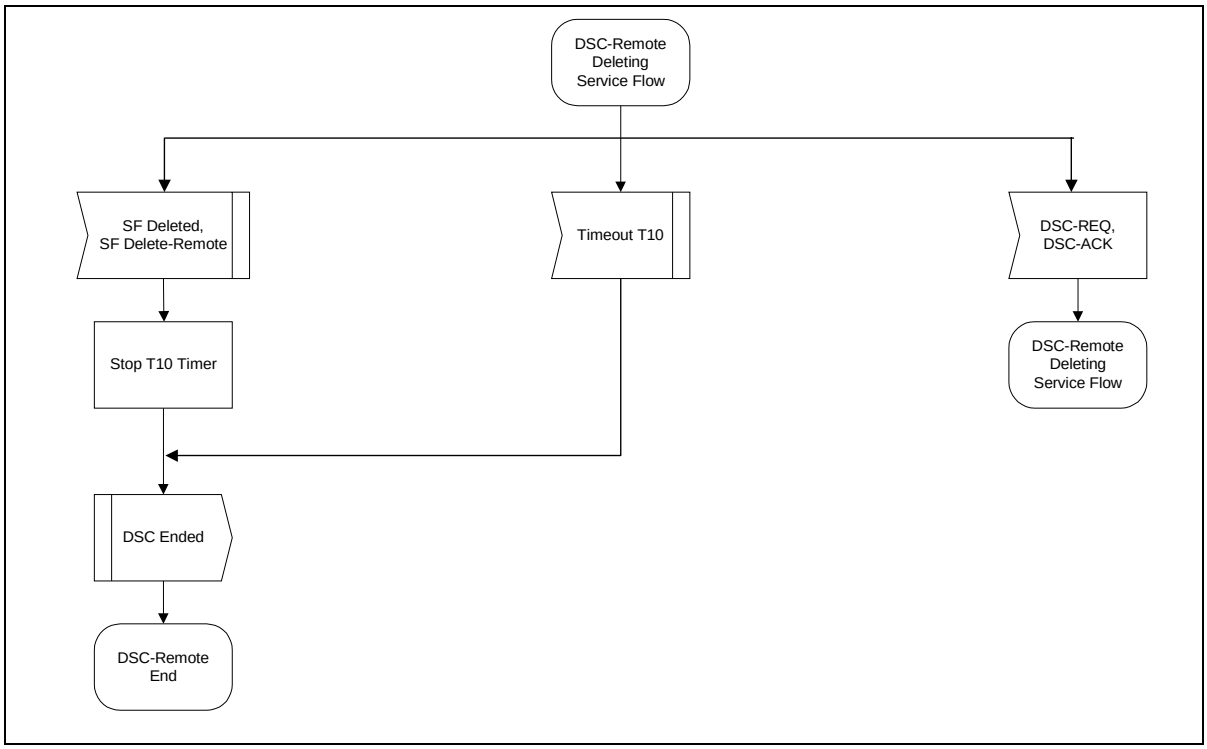


Figure 9-49. DSC—Remotely Initiated Transaction Deleting Service Flow State Flow Diagram

9.4.4 Dynamic Service Deletion

Any service flow can be deleted with the Dynamic Service Deletion (DSD) messages. When a Service Flow is deleted, all resources associated with it are released, including classifiers and PHS. However, if a Primary Service Flow of a CM is deleted, that CM is de-registered and MUST re-register. Also, if a Service Flow that was provisioned during registration is deleted, the provisioning information for that Service Flow is lost until the CM re-registers. However, the deletion of a provisioned Service Flow MUST NOT cause a CM to re-register. Therefore, care should be taken before deleting such Service Flows.

9.4.4.1 CM Initiated Dynamic Service Deletion

A CM wishing to delete an upstream and/or a downstream Service Flow generates a delete request to the CMTS using a Dynamic Service Deletion-Request message (DSD-REQ). The CMTS removes the Service Flow(s) and generates a response using a Dynamic Service Deletion-Response message (DSD-RSP). Only one upstream and/or one downstream Service Flow can be deleted per DSD-Request.

CM		CMTS
Service Flow(s) no longer needed		
Delete Service Flow(s)		
Send DSD-REQ	**---DSD-REQ-->**	**Receive DSD-REQ**
		Verify CM is Service Flow(s) 'owner'
		Delete Service Flow(s)
Receive DSD-RSP	**<--DSD-RSP---**	**Send DSD-RSP**

Figure 9-50. Dynamic Service Deletion Initiated from CM

9.4.4.2 CMTS Initiated Dynamic Service Deletion

A CMTS wishing to delete an upstream and/or a downstream dynamic Service Flow generates a delete request to the associated CM using a Dynamic Service Deletion-Request message (DSD-REQ). The CM removes the Service Flow(s) and generates a response using a Dynamic Service Deletion-Response message (DSD-RSP). Only one upstream and/or one downstream Service Flow can be deleted per DSD-Request.

CM		CMTS
		Service Flow(s) no longer needed
		Delete Service Flow(s)
		Determine associated CM for this Service Flow(s)
Receive DSD-REQ	**<---DSD-REQ---**	**Send DSD-REQ**
Delete Service Flow(s)		
Send DSD-RSP	**---DSD-RSP-->**	**Receive DSD-RSP**

Figure 9-51. Dynamic Service Deletion Initiated from CMTS

9.4.4.3 Dynamic Service Deletion State Transition Diagrams

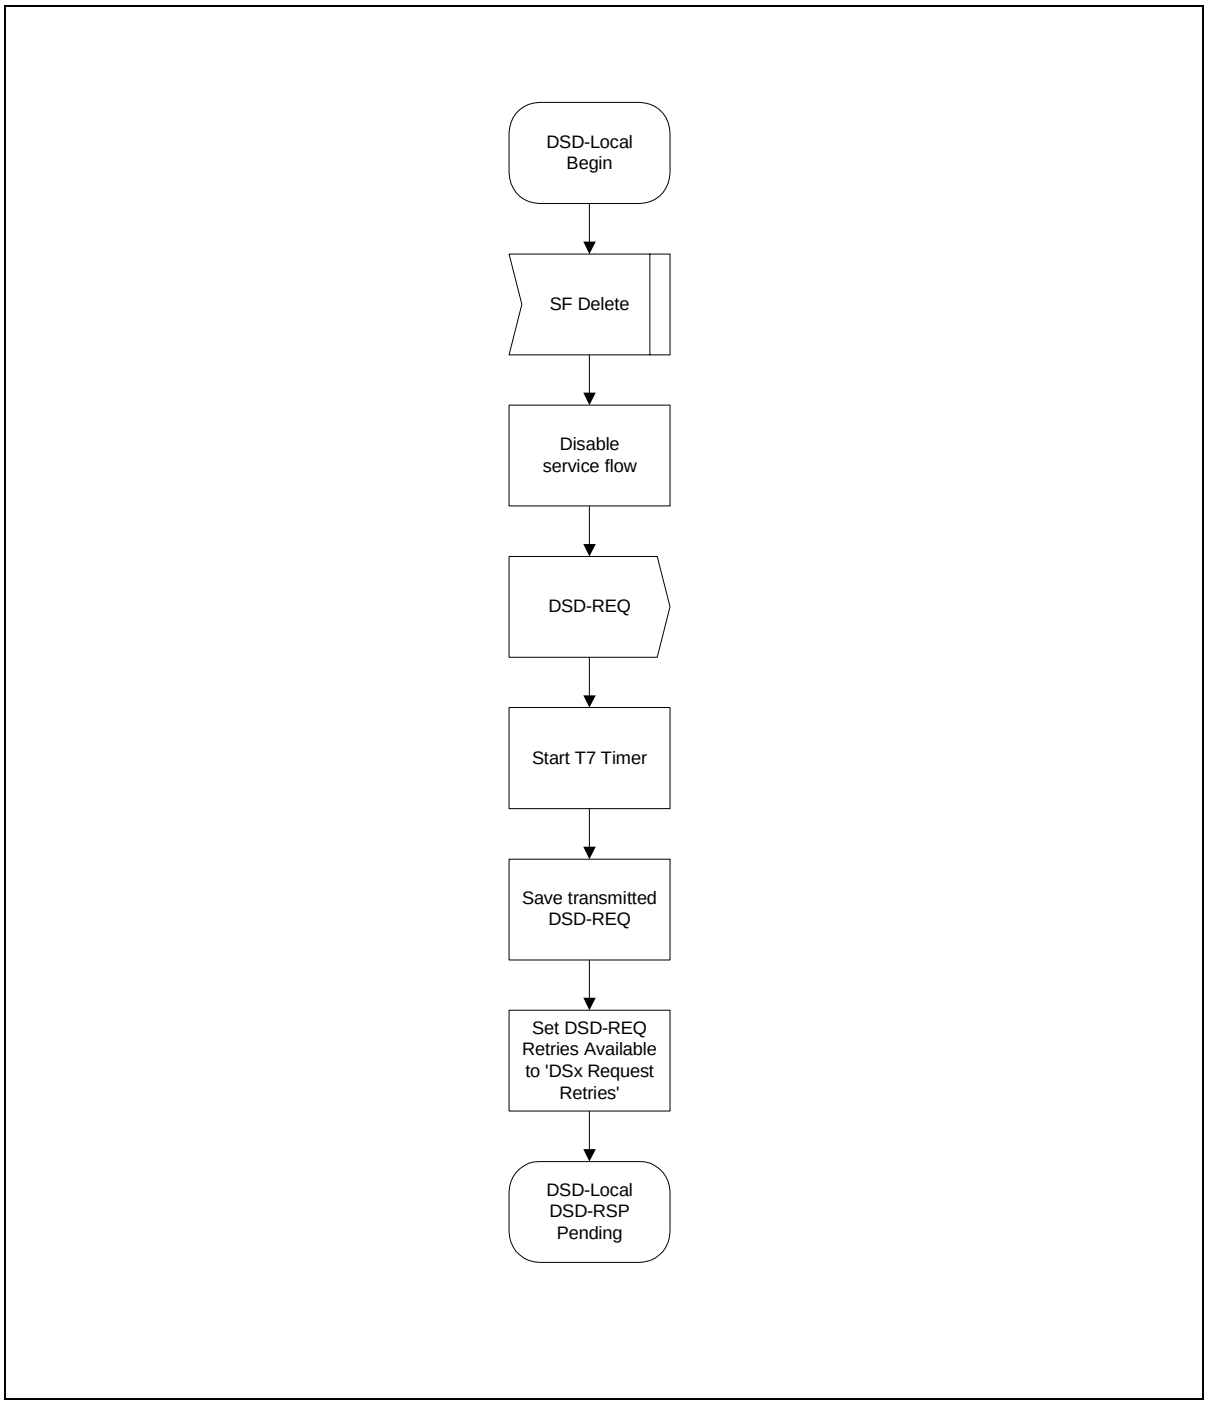


Figure 9-52. DSD—Locally Initiated Transaction Begin State Flow Diagram

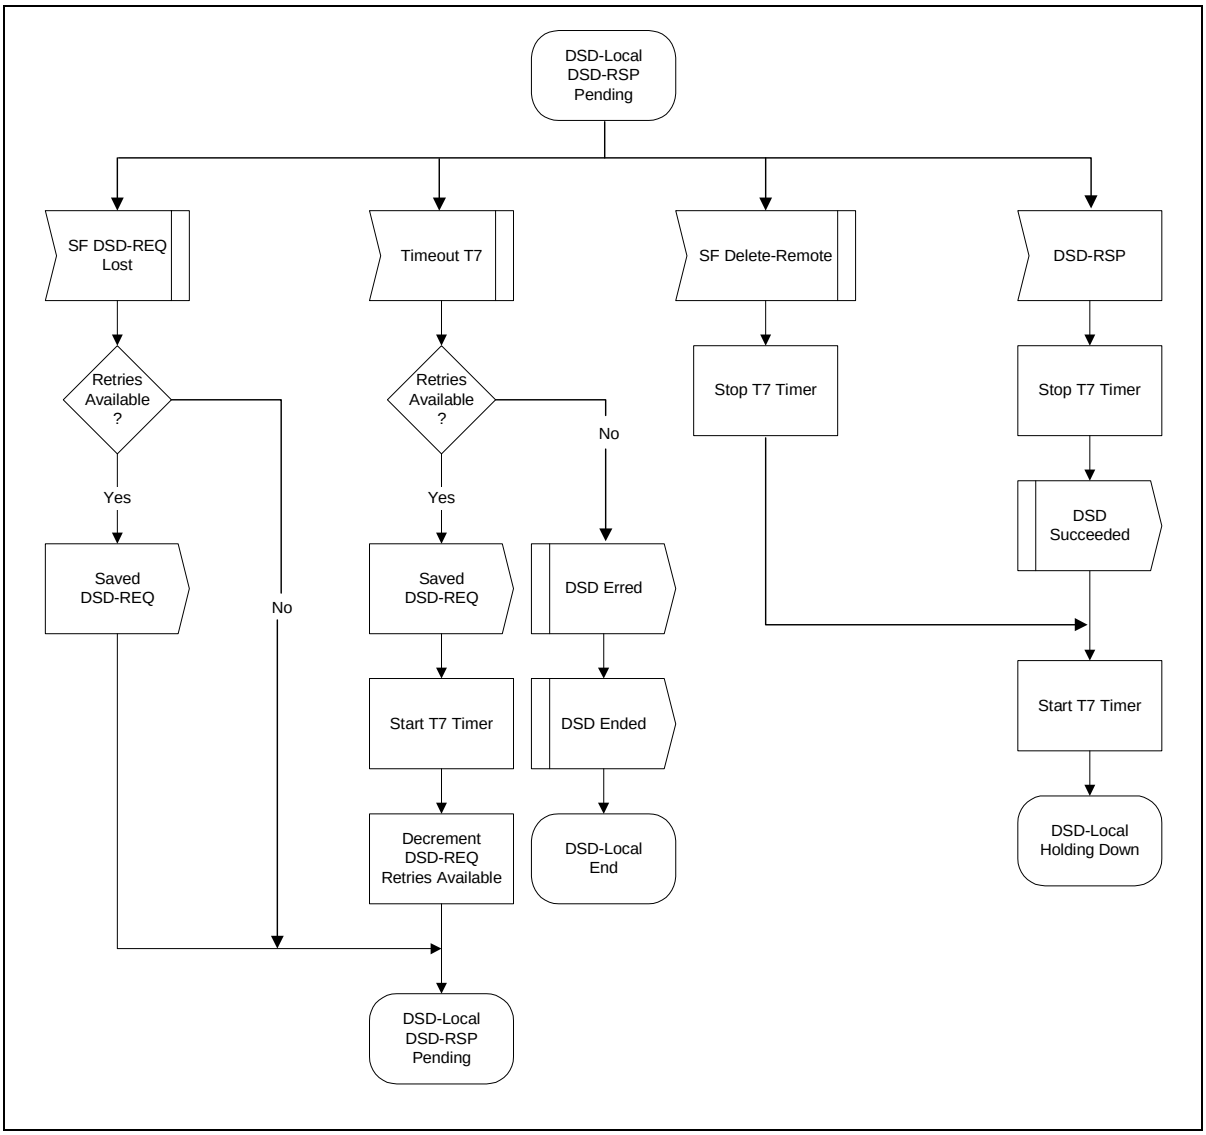


Figure 9-53. DSD—Locally Initiated Transaction DSD-RSP Pending State Flow Diagram

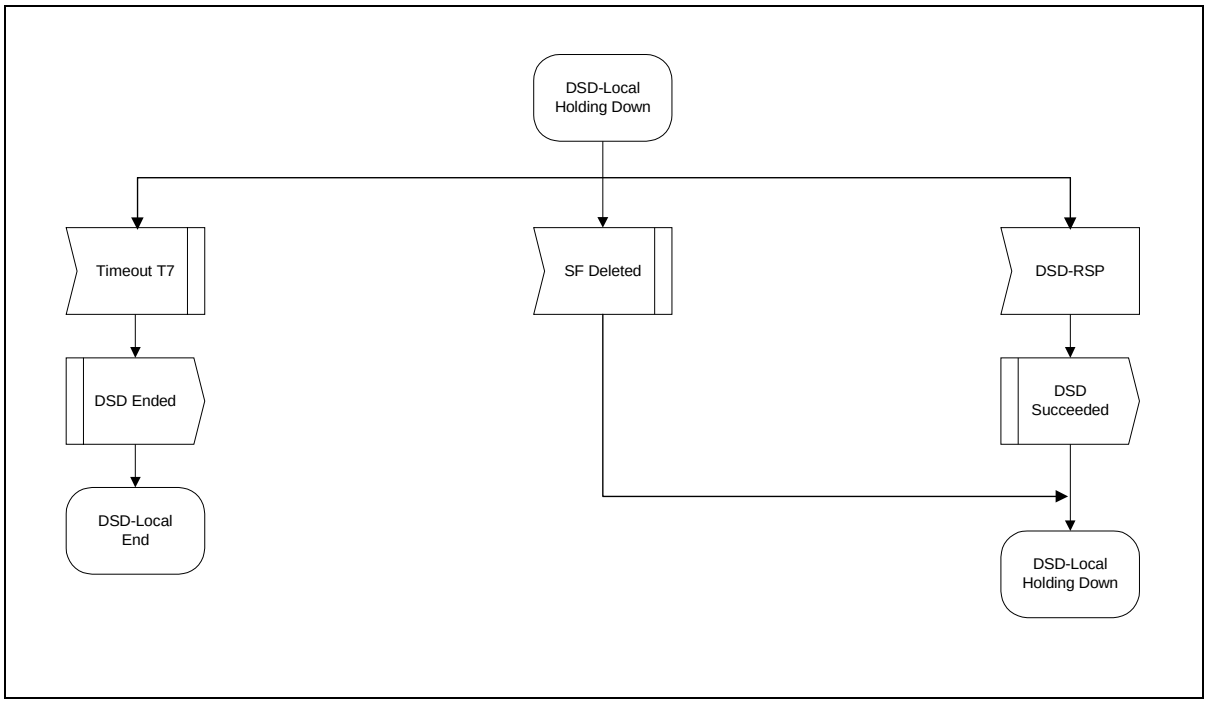


Figure 9-54. DSD—Locally Initiated Transaction Holding Down State Flow Diagram

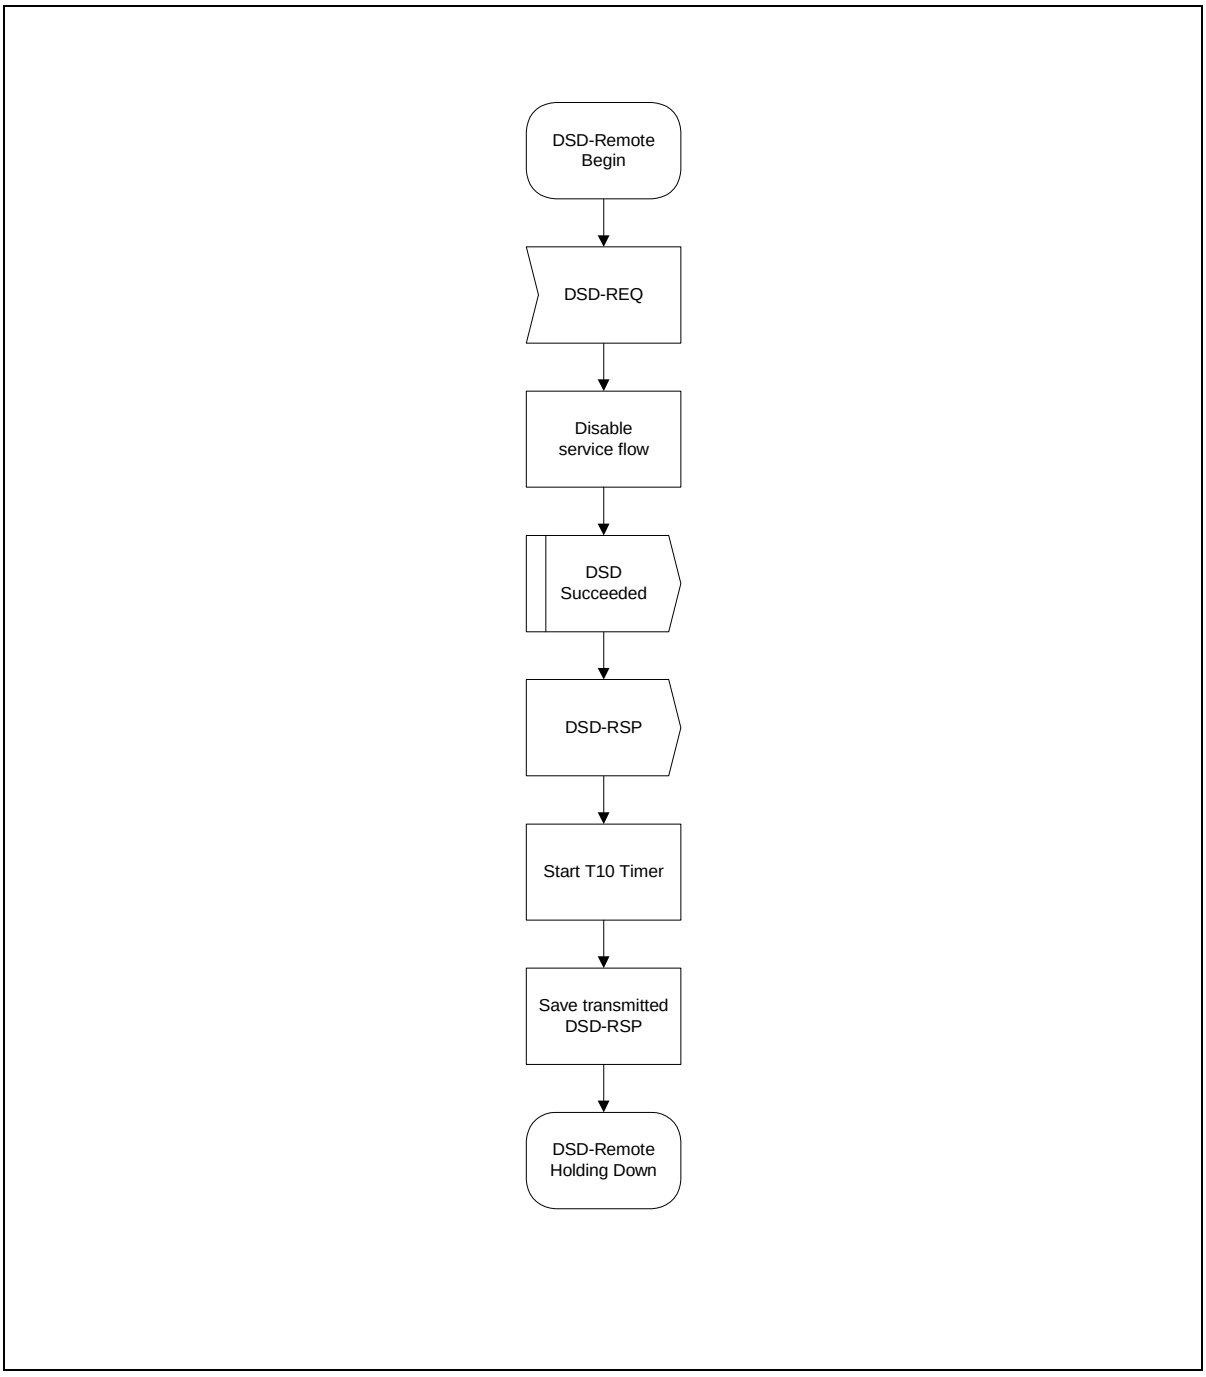


Figure 9-55. DSD—Remotely Initiated Transaction Begin State Flow Diagram

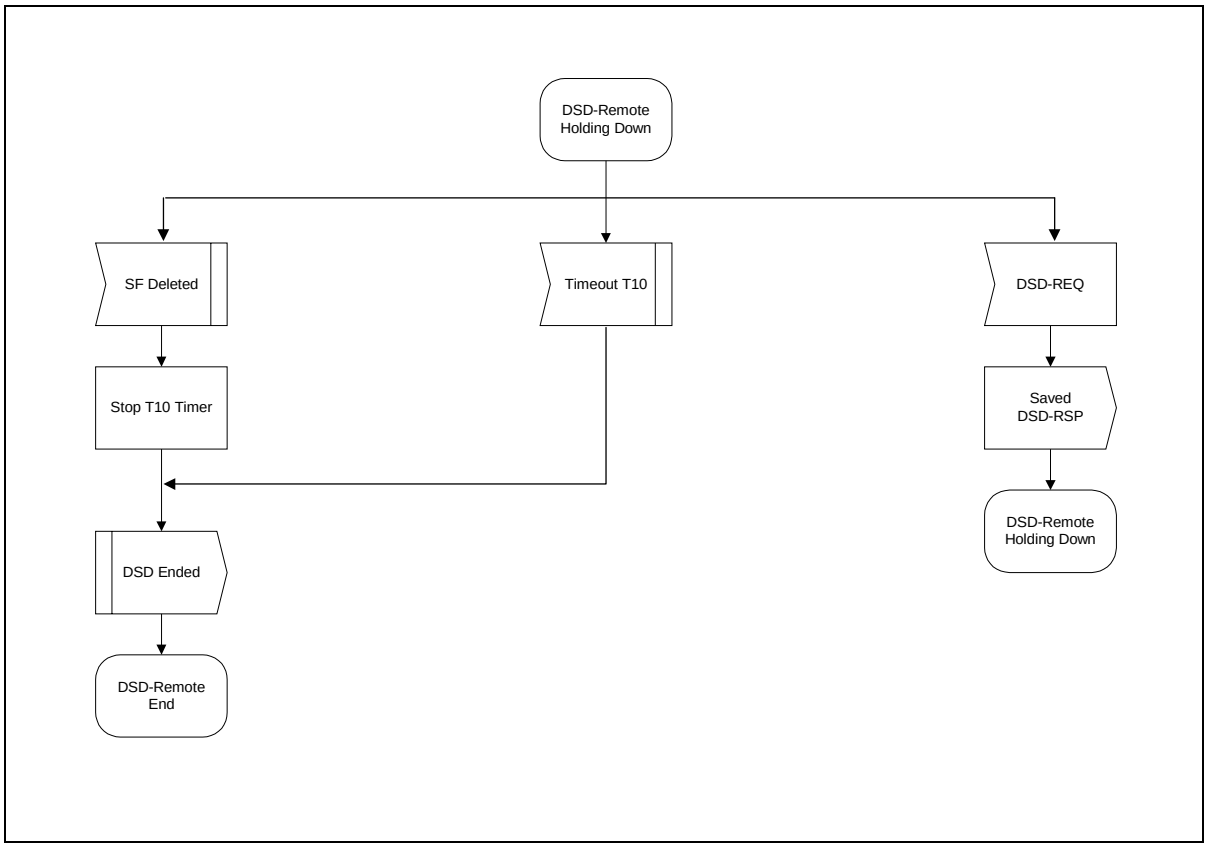


Figure 9-56. DSD—Remotely Initiated Transaction Holding Down State Flow Diagram

9.4.5 Dynamically Changing Downstream and/or Upstream Channels

9.4.5.1 DCC General Operation

At any time after registration, the CMTS MAY direct the CM to change its downstream and/or upstream channel. This may be done for traffic balancing, noise avoidance, or other reasons which are beyond the scope of this specification. Figure 9-58 and Figure 9-59 show the procedure that MUST be followed by the CMTS. Figure 9-62 shows the corresponding procedure that MUST be followed by a CM.

The DCC command can be used to change only the upstream frequency, only the downstream frequency, or both the upstream and downstream frequencies. When only the upstream or only the downstream frequency is changed, the change is typically within a MAC domain. When both the upstream and downstream frequencies are changed, the change may be within a MAC domain, or between MAC domains.

The Upstream Channel ID MUST be unique between the old and new channels. In this context, the old channel refers to the channel that the CM was on before the jump, and the new channel refers to the channel that the CM is on after the jump.

Upon synchronizing with the new upstream and/or downstream channel, the CM MUST use the technique specified in the DCC-REQ Initialization Technique TLV, if present, to determine if it should perform re-initialization, only ranging, or neither. If this TLV is not present in DCC-REQ, the CM MUST re-initialize its MAC on the new channel assignment. (Refer to Section 9.2). If the CM has been instructed to re-initialize, then the CMTS MUST NOT wait for a DCC-RSP to occur on the new channel.

If the CM is being moved within a MAC domain, then a re-initialization may not be required. If the CM is being moved between MAC domains, then a re-initialization may be required. Re-initializing, if requested, is done with the new upstream and channel assignments. It includes obtaining upstream parameters, establish IP connectivity, establish time of day, transfer operational parameters, register, and initialize baseline privacy. If re-initialization is performed, the CM MUST NOT send a DCC-RSP on the new channel.

The decision to re-range is based upon the CMTS's knowledge of any path diversity that may exist between the old and new channels, or if any of the fundamental parameters of the upstream or downstream channel such as symbol rate, modulation type, or minislot size have changed.

When DCC-REQ does not involve re-initialization or re-ranging, the design goal of the CM will typically be to minimize the disruption of traffic to the end user. To achieve this goal, a CM MAY choose to continue to use QoS resources (such as bandwidth grants) on its current channel after receiving a DCC-REQ and before actually executing the channel change. The CM might also need this time to flush internal queues or reset state machines prior to changing channels.

The CM MAY continue to use QoS resources on the old channel, including the transmission and reception of packets, after sending a DCC-RSP (depart) message and prior to the actual jump. The CM MAY use QoS resources on the new channel, including the transmission and reception of packets, after the jump and prior to sending a DCC-RSP (arrive) message. The CMTS MUST NOT use the DCC-RSP (depart) message to remove QoS resources on the old channel. The CMTS MUST NOT wait for a DCC-RSP (arrive) message on the new channel before allowing QoS resources to be used. This provision is to allow the Unsolicited Grant Service to be used on the old and new channel with a minimum amount of disruption when changing channels.

The CMTS MUST hold the QoS resources on the current channel until a time of T13 has passed after the last DCC-REQ that was sent, or until it can internally confirm the presence of the CM on the new channel assignment. The CM MUST execute the departure from the old channel before the expiry of T13. The CM MAY continue to use QoS resources on the old channel after responding with DCC-RSP and before the expiry of T13.

If the CM is commanded to perform initial or station maintenance or to use the channel directly, the destination CMTS MUST hold the QoS resources on the new channel until a time of T15 has passed after the last DCC-REQ was sent if the CM has not yet been detected on the new channel. If the CM is commanded to re-initialize the MAC, then QoS resources are not reserved on the destination CMTS, and T15 does not apply.

The T15 timer represents the maximum time period for the CM to complete the move to the destination CMTS, and is based on the TLV encodings (i.e., initialization technique TLV, UCD substitution TLV, and SYNC substitution TLV) included in the DCC-REQ message and the local configuration of the destination CMTS (UCD transmit interval, SYNC interval, etc.).

The destination CMTS SHOULD calculate and limit T15 based on internal policy according to the guidelines in Section 9.4.5.1.1.

If initialization technique of initial ranging is utilized and if the CM arrives after T15 has passed, attempting to use resources on the new channel that have been removed (ranging or requesting bandwidth on a SID that has been deleted), the CMTS MUST send a Ranging Abort to the CM in order to cause the DCC transaction to fail.

When a CM is moved between DS channels on different IP subnets using initialization techniques other than re-initialize the MAC, a network connectivity issue may occur because no DHCP process is indicated as part of the DCC operation. The CM MAY implement a vendor-specific feature to deal with this situation. The CMTS SHOULD take this issue into account when sending a DCC-REQ and SHOULD direct the CM to use the appropriate initialization technique TLV to ensure no IP connectivity loss as a result of DCC.

Once the CM changes channels, all previous outstanding bandwidth requests made via the Request IE or Request/Data IE are invalidated, and the CM MUST re-request bandwidth on the new channel. In the case of Unsolicited Grant Service in the upstream, the grants are implicit with the QoS reservations, and do not need to be re-requested.

9.4.5.1.1 Derivation of T15 Timer

The maximum value noted for the T15 timer denotes the maximum amount of time that the CMTS should reserve resources on the new channel. This value is not to be used to represent acceptable performance.

The equation below describes the method for calculating the value of T15.

T15 = CmJumpTime + CmRxTargetUcd + CmRxDsSync + CmtsRxRngReq

Each of the variables in the equation calculating the T15 timer is explained in the table below.

Variable	Explanation and Value
CmJumpTime	This is the CM's indication to the CMTS of when it intends to start the jump and how long it will take to jump. For a downstream change, it includes the time for the CM to synchronize to the downstream parameters on the destination channel, such as QAM symbol timing, FEC framing, and MPEG framing. It incorporates CM housecleaning on the old channel. It also incorporates one T11 period for the CM to process and receive the DCC-REQ message. This optional value is computed by CM and returned in DCC-RSP (depart). If CM does not specify the Jump Time TLVs, then the destination CMTS assumes that the value is 1.3 seconds. This recognizes the fact that the CM may continue to use the old channel until the expiry of the T13 timer. If CM specifies the Jump Time TLVs, then the destination CMTS uses the specified value.
CmRxTargetUcd	This variable represents the time for the CM to acquire UCD parameters for the target upstream channel. The value of this variable is two CMTS UCD timer periods.
CmRxDsSync	This variable represents the time for the CM to acquire a downstream SYNC message. The value of this variable is two CMTS SYNC timer periods.
CmtsRxRngReq	This variable represents the time for the CM to receive and use a ranging opportunity, and for the CMTS to receive and process the RNG-REQ. For the initialization technique of use directly, this value is two times the CMTS time period between unicast ranging opportunities plus 20 - 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time. For the initialization technique of station maintenance, this value is two times the CMTS time period between unicast ranging opportunities plus 20 - 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time. For the initialization technique of initial maintenance, this value is 30 seconds. Because the variables involved in initial maintenance are not strictly under the control of the CMTS, the computation of this factor is uncertain.

The maximum value assigned to the T15 timer denotes the maximum amount of time that the CMTS should reserve resources on the new channel. The minimum value of the T15 timer is two seconds; this was derived by doubling the value of the T13 timer. The maximum value of the T15 timer is 35 seconds.

9.4.5.2 DCC Exception Conditions

If a CM issues a DSA-REQ or DSC-REQ for more resources, and the CMTS needs to do a DCC to obtain those resources, the CMTS will reject the DSA or DSC command without allocating any resources to the CM. The CMTS includes a confirmation code of "reject-temporary-DCC" (refer to Appendix C.1.3.1) in the DSC-RSP message to indicate that the new resources will not be available until a DCC is received. The CMTS will then follow the DSA or DSC transaction with a DCC transaction.

After the CM jumps to a new channel and completes the DCC transaction, the CM retries the DSA or DSC command. If the CM has not changed channels after the expiry of T14, as measured from the time that the CM received DSA-RSP or DSC-RSP from the CMTS, then the CM MAY retry the resource request.

If the CMTS needs to change channels in order to satisfy a resource request other than a CM initiated DSA or DSC command, then the CMTS should execute the DCC command first, and then issue a DSA or DSC command.

If a CMTS does a DCC with re-initialize, the config file could cause the CM to come back to the original channel. This would cause an infinite loop. To prevent this, if the provisioning system default is to specify the upstream channel ID and/or the downstream frequency, then the CMTS SHOULD NOT use DCC-REQ with the re-initialize option.

The CMTS MUST NOT issue a DCC command if the CMTS has previously issued a DSA, or DSC command, and that command is still outstanding. The CMTS MUST NOT issue a DCC command if the CMTS is still waiting for a DSA-ACK or DSC-ACK from a previous CM initiated DSA-REQ or DSC-REQ command.

The CMTS MUST NOT issue a DSA or DSC command if the CMTS has previously issued a DCC command, and that command is still outstanding.

If the CMTS issues a DCC-REQ command and the CM simultaneously issues a DSA-REQ or DSC-REQ then the CMTS command takes priority. The CMTS responds with a confirmation code of "reject-temporary" (refer to Appendix C.1.3.1). The CM proceeds with executing the DCC command.

If the CM is unable to achieve communications with a CMTS on the new channel(s), it MUST return to the previous channel(s) and re-initialize its MAC. The previous channel assignment represents a known good operating point which should speed up the re-initialization process. Also, returning to the previous channel provides a more robust operational environment for the CMTS to find a CM that fails to connect on the new channel(s).

If the CMTS sends a DCC-REQ and does not receive a DCC-RSP within time T11, it MUST retransmit the DCC-REQ up to a maximum of "DCC-REQ Retries" (Appendix B) before declaring the transaction a failure. Note that if the DCC-RSP was lost in transit and the CMTS retries the DCC-REQ, the CM may have already changed downstream channels.

If the CM sends a DCC-RSP on the new channel and does not receive a DCC-ACK from the CMTS within time T12, it MUST retry the DCC-RSP up to a maximum of "DCC-RSP Retries" (Appendix B).

If the CM receives a DCC-REQ with the Upstream Channel ID TLV, if present, equal to the current Upstream Channel ID, and the Downstream Frequency TLV, if present, is equal to the current downstream frequency, then the CM MUST consider the DCC-REQ as a redundant command. The remaining DCC-REQ TLV parameters MUST NOT be executed, and the CM MUST return a DCC-RSP, with a confirmation code of "reject-already-there", to the CMTS (refer to Appendix C.4.1).

9.4.5.3 DCC Performance

The purpose of a DCC is to move the CM to a new upstream and/or downstream channel with little interruption of service. There are many factors that affect the performance of a DCC transaction including CM purging, initialization technique, and the number of TLV hints given by the current CMTS in the DCC-REQ message. Each of these factors is individually discussed in the table below.

The DCC transaction is defined from the perspective of both the CM and the CMTS for the discussion on performance in the following table. From the perspective of the CM, the DCC transaction begins when the CM receives the DCC-REQ message from the CMTS and completes when the CM receives the DCC-ACK message from the CMTS. From the perspective of the CMTS, the DCC transaction begins when the CMTS sends the DCC-REQ message to the CM and completes when the CMTS receives the DCC-RSP (arrive) message from the CM.

TLV Type	Value	Explanation
Initialization Technique	Absent or 0 Reinitialize MAC	There are no performance requirements in this case. The CM will arrive on the destination CMTS after initialization occurs.
	1 Initial Maintenance	There are low performance expectations in this case because many factors affect the performance, such as collisions and ranging backoff. The CM should arrive on the destination CMTS as quickly as possible.
	2 Station Maintenance	The DCC transaction SHOULD complete within 1.5 seconds after the start of jump if the UCD substitution TLV and the downstream parameter TLVs are supplied. The DCC transaction SHOULD complete within the sum of CM jump time, two UCD intervals, and two ranging intervals if the current CMTS supplies no TLV hints in the DCC-REQ message.
	3 Initial or Station Maintenance	The CMTS does not know which ranging technique the CM will utilize. The CM should arrive on the destination CMTS as quickly as possible.
	4 Use Channel Directly	The DCC transaction SHOULD complete within one second after the start of jump if the UCD substitution TLV and the downstream parameter TLVs are supplied. The DCC transaction SHOULD occur within the sum of CM jump time and two UCD intervals if the current CMTS supplies no TLV hints in the DCC-REQ message.
DS Parameter		The CMTS SHOULD include the downstream parameter TLVs for station maintenance and use directly initialization techniques that are expected to occur quickly.
UCD Substitution		The CMTS SHOULD include the UCD substitution TLV for station maintenance and use directly initialization techniques that are expected to occur quickly.
SYNC Substitution		The CMTS SHOULD include the SYNC substitution TLV for station maintenance and use directly initialization techniques that are expected to occur quickly.
CM Jump Time		The length of jump TLV SHOULD be less than one second for downstream channel changes that include the downstream parameter TLVs or for upstream only channel changes.

When the DCC-REQ does not contain UCD Substitution TLVs and/or specifies an Initialization Technique of Initial Maintenance, Station Maintenance, or use directly, the destination CMTS SHOULD increase the probability that the CM will arrive quickly by using the CM Jump Time TLVs specified in the DCC-RSP (depart) to adjust the transmission of UCDs and ranging opportunities such that they coincide with the time when CM has estimated that it will arrive, and SHOULD increase the frequency of UCDs and/or ranging opportunities during this period.

9.4.5.4 Near-Seamless Channel Change

When the CMTS wishes to add new QoS reservations to a CM, it may be necessary to move that CM to a new upstream and/or downstream to achieve that goal. During that changing of channels, it is desirable to provide the minimum of interruption to existing QoS services such as voice over IP or video streaming sessions. This near-seamless channel change is the primary design goal of the DCC command. The CMTS MAY support a near-seamless channel change. The CM MAY support a near-seamless channel change.

The actions below are recommended operating procedures to implement a near-seamless channel change. The list assumes both the upstream and downstream channels are changing. A subset of the list would apply if only the upstream or downstream channel changed.

To support a near-seamless channel change, the following conditions should apply in the network:

- The physical layer parameters for the new upstream and downstream channels should not change with the old upstream and downstream channels. Note that a change in downstream parameters could invalidate the ranging parameters.
- The ranging parameters should not change between the old and new channels. This may require symmetrical cabling and plant conditions which are external to the CMTS.
- The CMTS should use the same time stamp and SYNC mechanism for all downstream channels.
- IP routing should be configured so that the CM and its attached CPEs can continue to use their existing IP addresses. This will avoid disruption to RTP sessions or other in progress applications.

To achieve a near-seamless channel change, the CMTS:

- SHOULD duplicate all the relevant QoS reservations for the CM on the old and new channel assignments before initiating a DCC-REQ.
- SHOULD duplicate downstream packet flow for the CM on the old and new channel assignments before initiating a DCC-REQ (for downstream channel changes).
- SHOULD transmit MAP messages for the new upstream channel on the old downstream channel for at least the duration of T13, if the old and new downstream channels share the same timestamp. (Note that if the CM cannot cache MAPs for the new upstream while on the old downstream channel, then the channel change delay will be increased by the amount of time into the future that MAPs are generated. Thus, the CMTS SHOULD refrain from scheduling MAPs farther into the future than it needs to.)
- SHOULD specify the downstream and upstream parameters of the new channels prior to the CM jumping.
- SHOULD specify to not wait for a SYNC message on the new channel.
- SHOULD specify to skip initialization (as defined in Section 9.2).
- SHOULD specify to skip initial maintenance and station maintenance.
- SHOULD manage service flow substitutions between old and new SIDs, SAID, Service Flow IDs, and Unsolicited Grant Time Reference as required. Service Class Names SHOULD remain the same between the old and new channel(s).

To achieve a near-seamless channel change, the CM:

- SHOULD reply with estimates for CM Jump Time in the DCC-RSP message.
- SHOULD listen for and cache MAP messages on the old downstream that apply to the new upstream. This SHOULD be done during time T13.
- SHOULD use the downstream parameters and the UCD in its cache from the DCC command to force a quicker PHY convergence when jumping.

- SHOULD NOT wait for a SYNC message after PHY convergence and before transmitting, if the CMTS permits the CM to do so.
- SHOULD use the cached MAPS, if available, to allow a quicker start-up time.
- SHOULD minimize the disruption of traffic in either direction by allowing traffic to continue to flow in both directions up to the moment prior to the jump and then immediately after resynchronization to the new channel(s) has happened.
- SHOULD queue incoming data packets that arrive during the jump, and transmit them after the jump
- SHOULD discard VoIP packets after the jump that have caused the upstream Unsolicited Grant Service queue to exceed its limit, but no more than necessary.

Applications that are running over the DOCSIS path should be able to cope with the loss of packets that may occur during the time that the CM changes channels.

9.4.5.5 Example Operation

9.4.5.5.1 Example Signaling

Figure 9-57 shows an example of the use of DCC and its relation to the other DOCSIS MAC messages. In particular, this example describes a scenario where the CM attempts to allocated new resources with a DSA message. The CMTS temporarily rejects the request, tells the CM to change channels, and then the CM re-requests the resources. This example (not including all exception conditions) is described below. Refer to Section 9.2 for more detail.

a) An event occurs, such as the CM issuing a DSA-REQ message.

b) The CMTS decides that it needs to change channels in order to service this resource request. The CMTS responds with a DSA-RSP message which includes a confirmation code of “reject-temporary-DCC” (refer to Appendix C.1.3.1) in the DSC-RSP message to indicate that the new resources are not available until a DCC is received. The CMTS now rejects any further DSA or DSC messages until the DCC command is executed.

c) The CMTS initiates QoS reservations on the new upstream and/or downstream channels. The QoS reservations include the new resource assignment along with all the current resource assignments assigned to the CM. In this example, both the upstream and downstream channels are changed.

d) To facilitate a near-seamless channel change, since the CMTS is not sure exactly when the CM will switch channels, the CMTS duplicates the downstream packet flow on the old and new downstream channels.

e) The CMTS issues a DCC-REQ command to the CM.

f) The CM sends a DCC-RSP (depart). The CM then cleans up its queues and state machines as appropriate and changes channels.

g) If there was a downstream channel change, the CM synchronizes to the QAM symbol timing, synchronizes the FEC framing, and synchronizes with the MPEG framing.

h) If the CM has been instructed to re-initialization, it does so with the new upstream and/or downstream channel assignment. The CM exits from the flow of events described here, and enters the flow of events described in Section 9.2 starting with the recognition of a downstream SYNC message.

i) The CM searches for a UCD message unless it has been supplied with a copy.

j) The CM waits for a downstream SYNC message unless it has been instructed not to wait for one.

k) The CM collects MAP messages unless it already has them available in its cache.

l) The CM performs initial maintenance and station maintenance unless it has been instructed to skip them.

m) The CM resumes normal data transmission with its new resource assignment.

n) The CM sends a DCC-RSP (arrive) message to the CMTS.

o) The CMTS responds with a DCC-ACK.

p) The CMTS removes the QoS reservations from the old channels. If the downstream packet flow was duplicated, the packet duplication would also be removed on the old downstream channel.

q) The CM re-issues its DSA-REQ command.

r) The CMTS reserves the requested resources and responds with a DSA-RSP.

s) The CM finishes with a DSA-ACK.

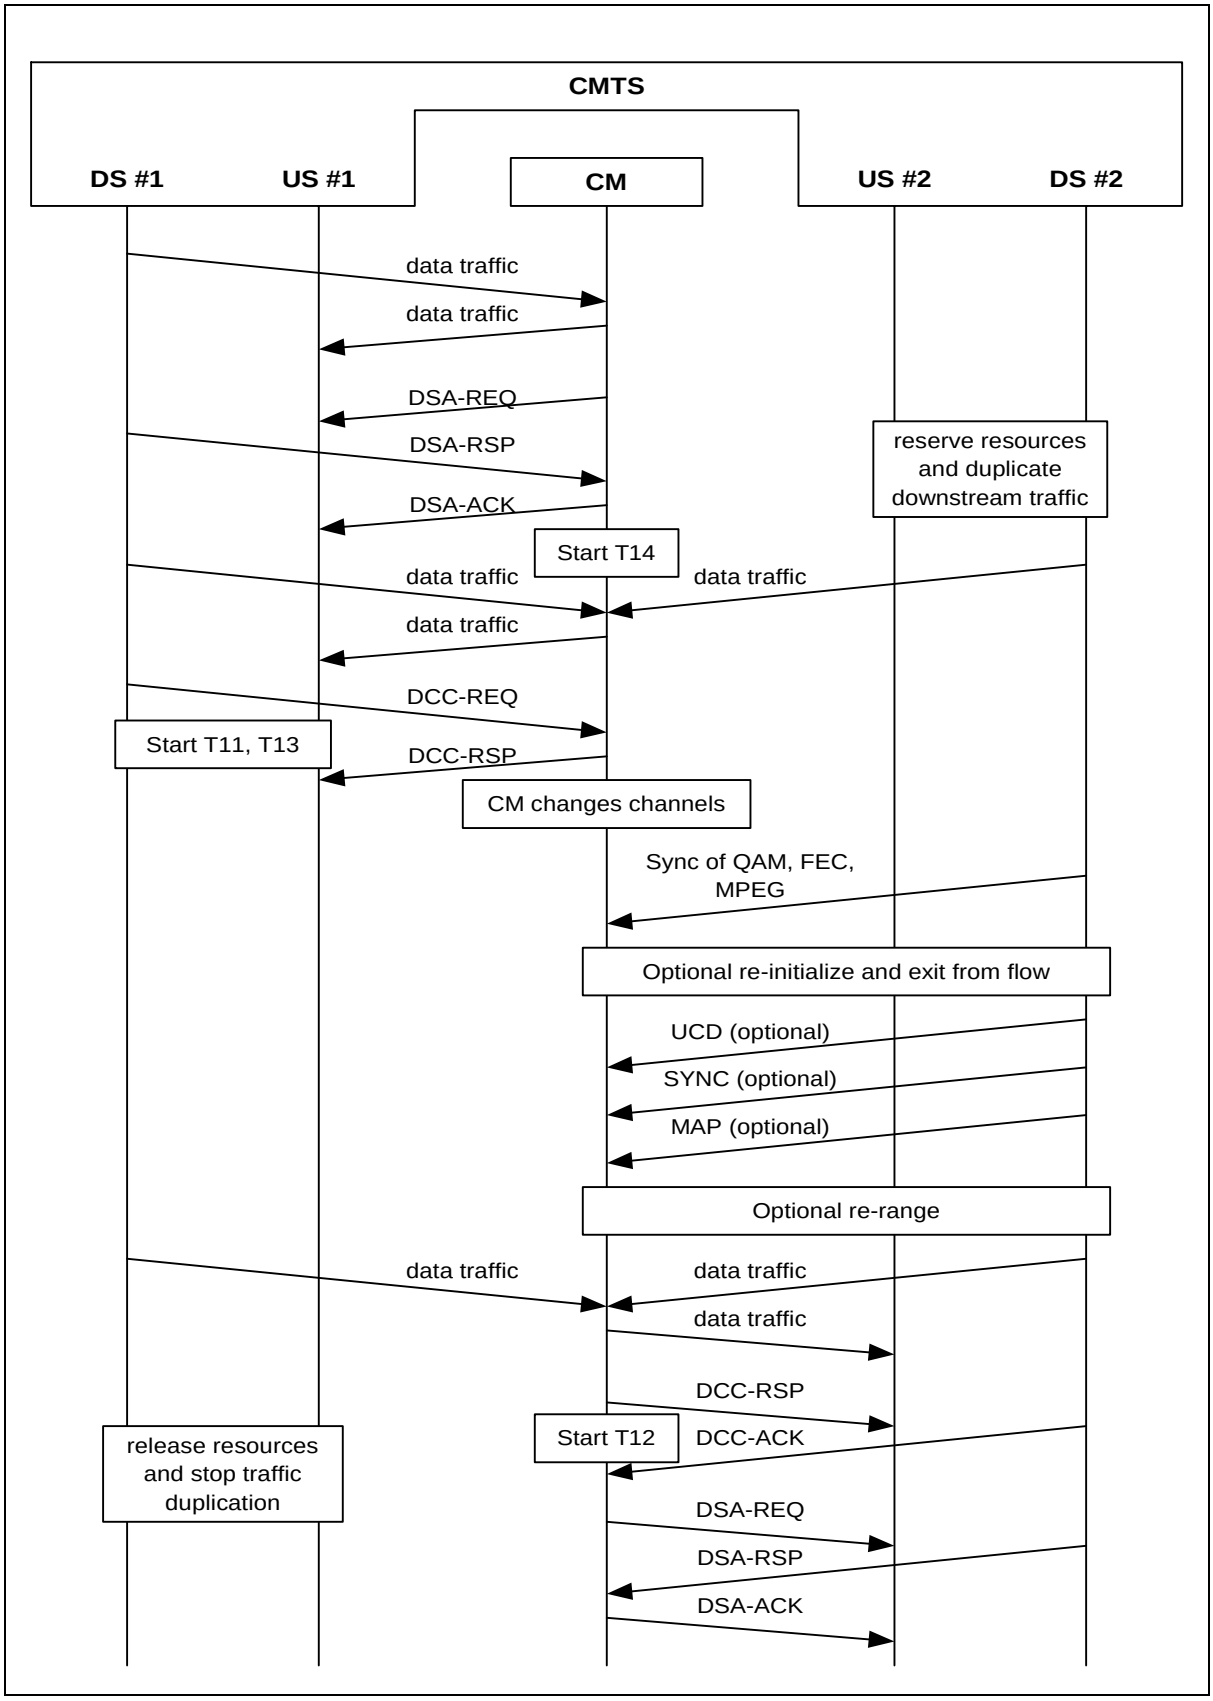


Figure 9-57. DCC Example Operational Flow

The states for the old and new CMTSes are shown as separate flow diagrams, since the old and new CMTS may be different. If the CMTSes are the same (e.g., the same MAC domain), then the CMTS will need to run both sets of state machines concurrently.

The flow diagrams show points where explicit signaling between the old and new CMTS is desirable, especially for near-seamless operation. The mechanism for this signaling is beyond the scope of this specification.

Note that the flow diagrams for both old and new CMTSs have been carefully crafted to handle many error conditions, such as:

- If the CM does not respond to the DCC-REQ (or responds with a reject conf code) and does not move, then it will be allowed to remain on the old channel. Resources on the new channel will be released (old CMTS signals DCC aborted to the new CMTS).
- If the CM DCC-RSP (depart) is lost, but the CM moves and arrives on the new CMTS, the new CMTS will signal that the CM has arrived to the old CMTS, allowing it to remove resources.
- If the CM DCC-RSP (depart) is received and the CM DCC-RSP (arrive) is lost, but the new CMTS otherwise detects the presence of the CM, the DCC transaction is considered successful, and the CM is allowed to remain on the new channel.
- If the CM DCC-RSP (depart) and (arrive) are lost, but the new CMTS otherwise detects the presence of the CM, the new CMTS will signal that the CM has arrived to the old CMTS, allowing it to remove resources, and the CM is allowed to remain on the new channel.
- If the CM DCC-RSP (depart) is received, but the CM never arrives, the new CMTS will detect this and remove resources after T15 expires.
- If the CM DCC-RSP (depart) is lost and the CM never arrives, the old CMTS will signal DCC aborted to the new CMTS, allowing it to remove resources. The old CMTS will use a different mechanism outside the scope of the DCC flow diagrams (such as ranging timeout) to remove resources on the old channels.
- If the CMTS DCC-ACK is lost and the DCC-RSP retry counter is expired, the CM will log an error and continue to the operational state.

There is a race condition that is not addressed in the flow diagrams; if the CM DCC-RSP (depart) is lost, the old CMTS will signal DCC aborted to the new CMTS. If the CM is in the process of moving, but has not yet arrived, the new CMTS will remove resources. This will prevent the CM from arriving successfully, unless it is able to complete the jump and arrive in approximately 1.2 seconds (3 retries of the DCC-REQ).

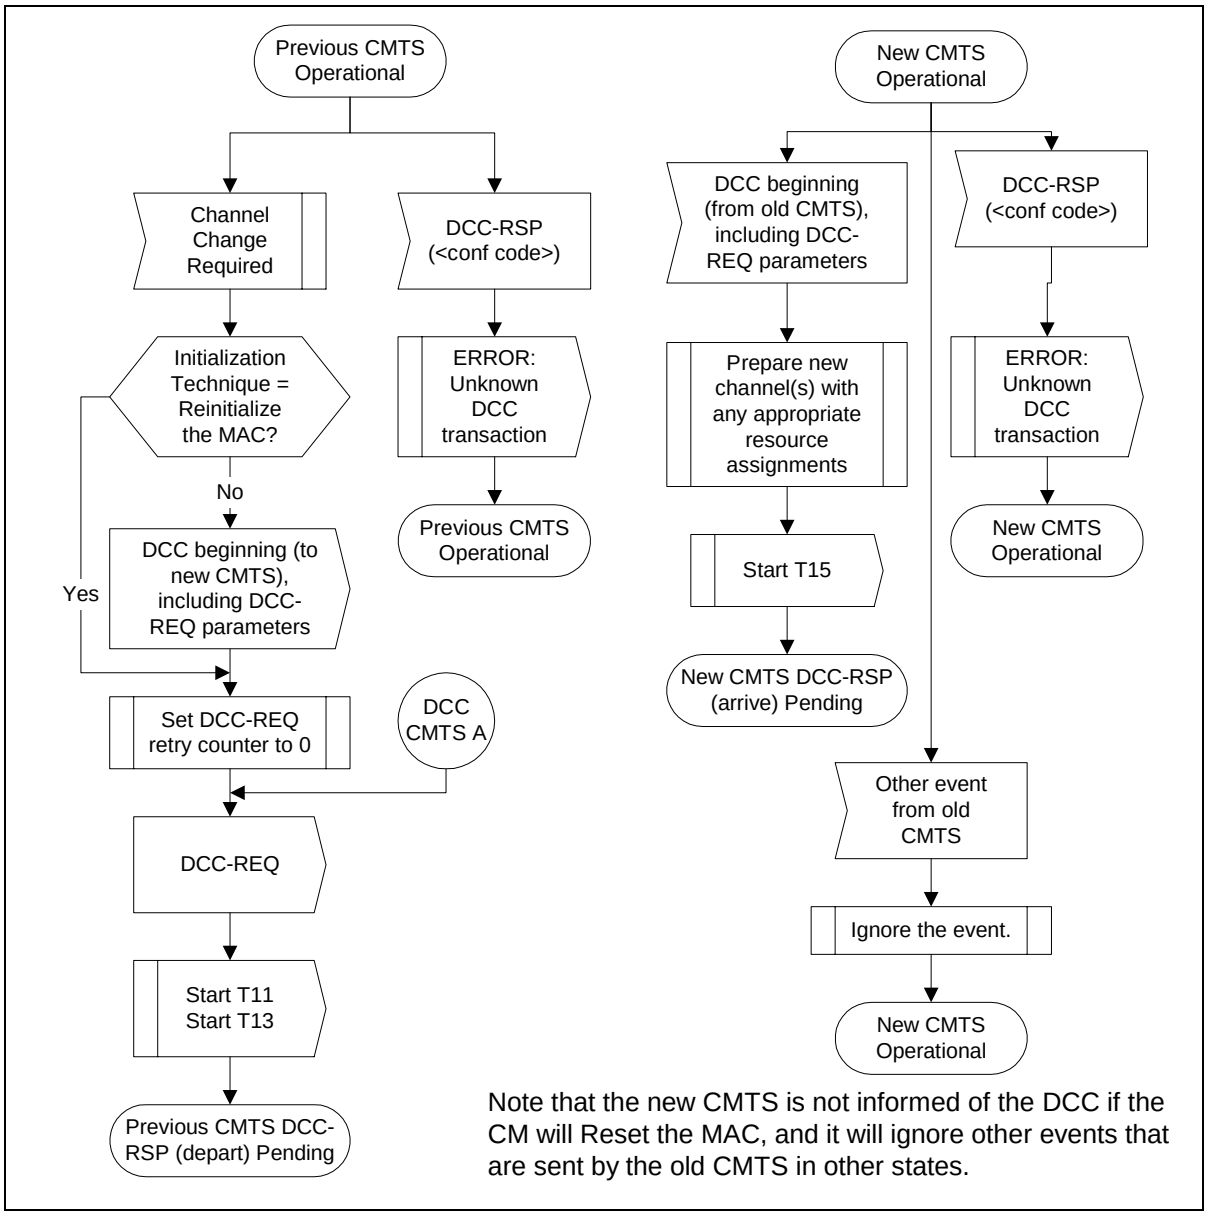


Figure 9-58. Dynamically Changing Channels: CMTS View Part 1

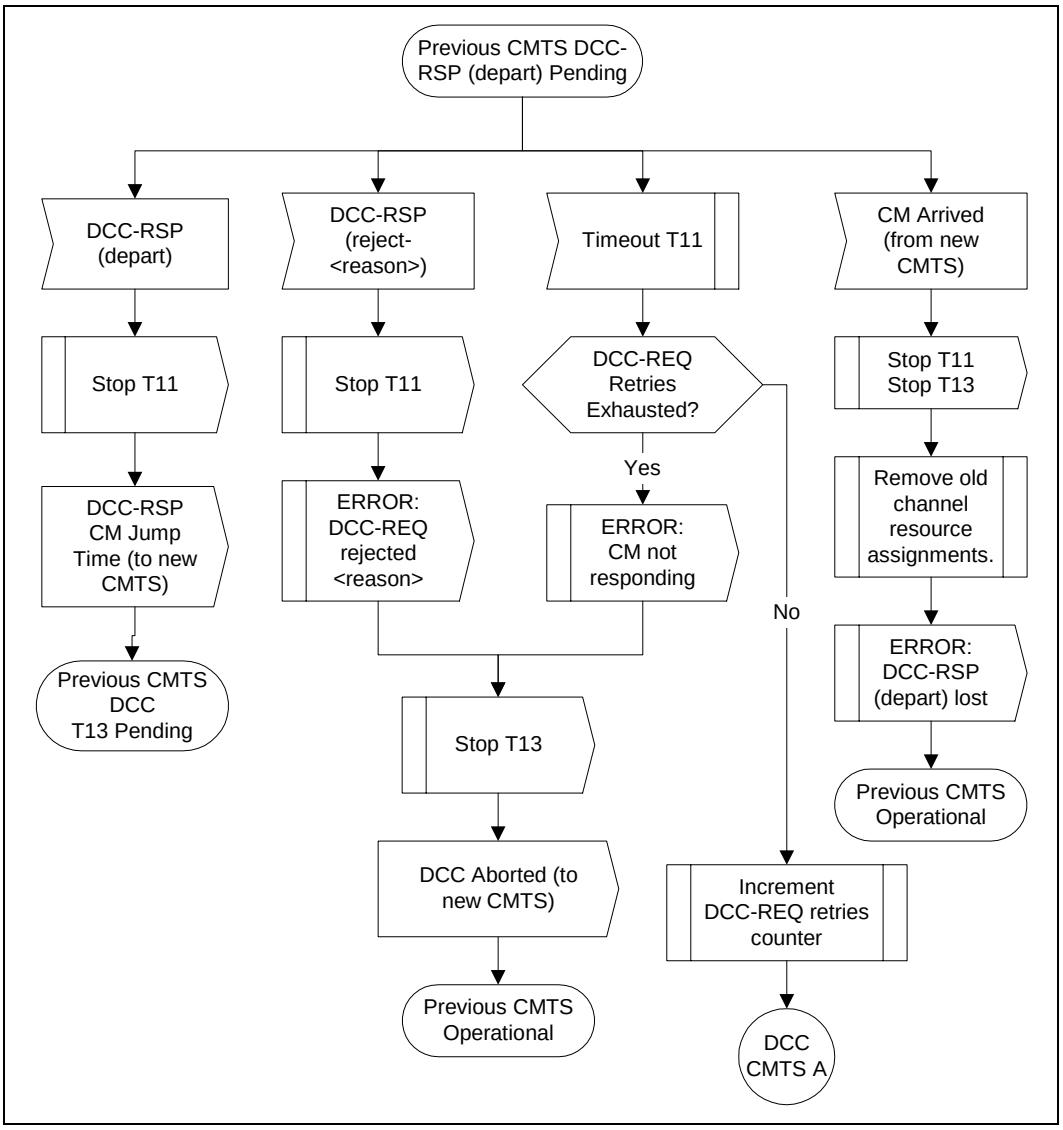


Figure 9-59. Dynamically Changing Channels: CMTS View Part 2

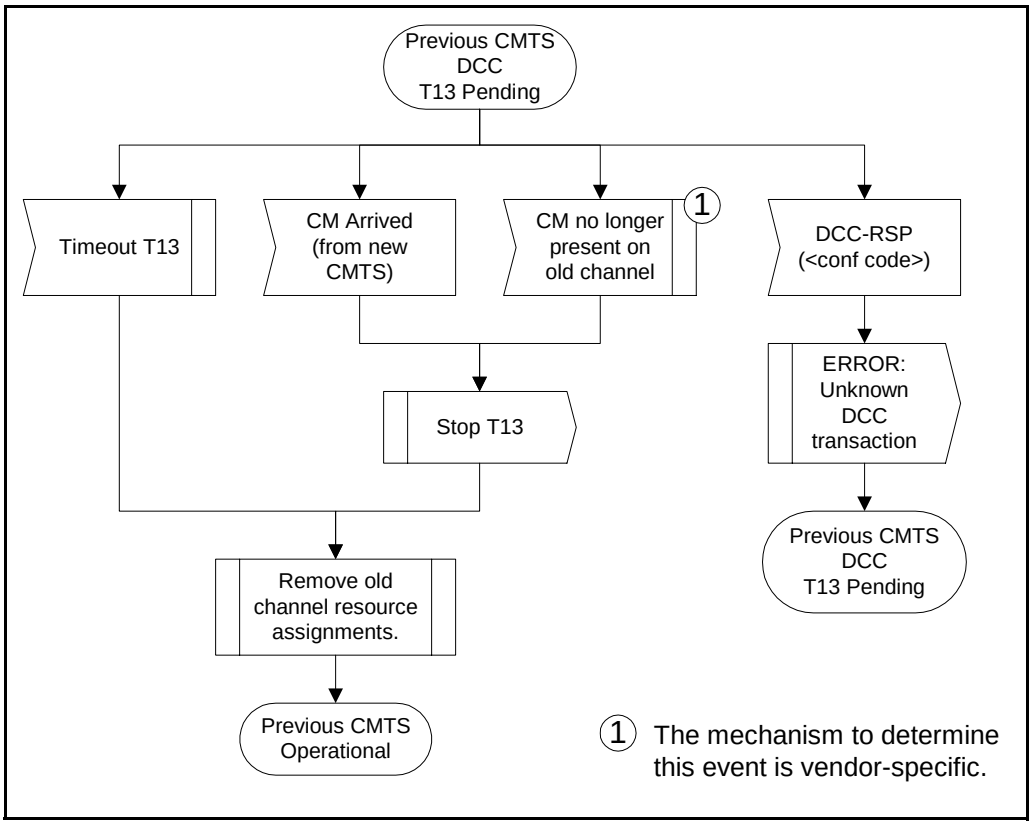


Figure 9-60. Dynamically Changing Channels: CMTS View Part 3

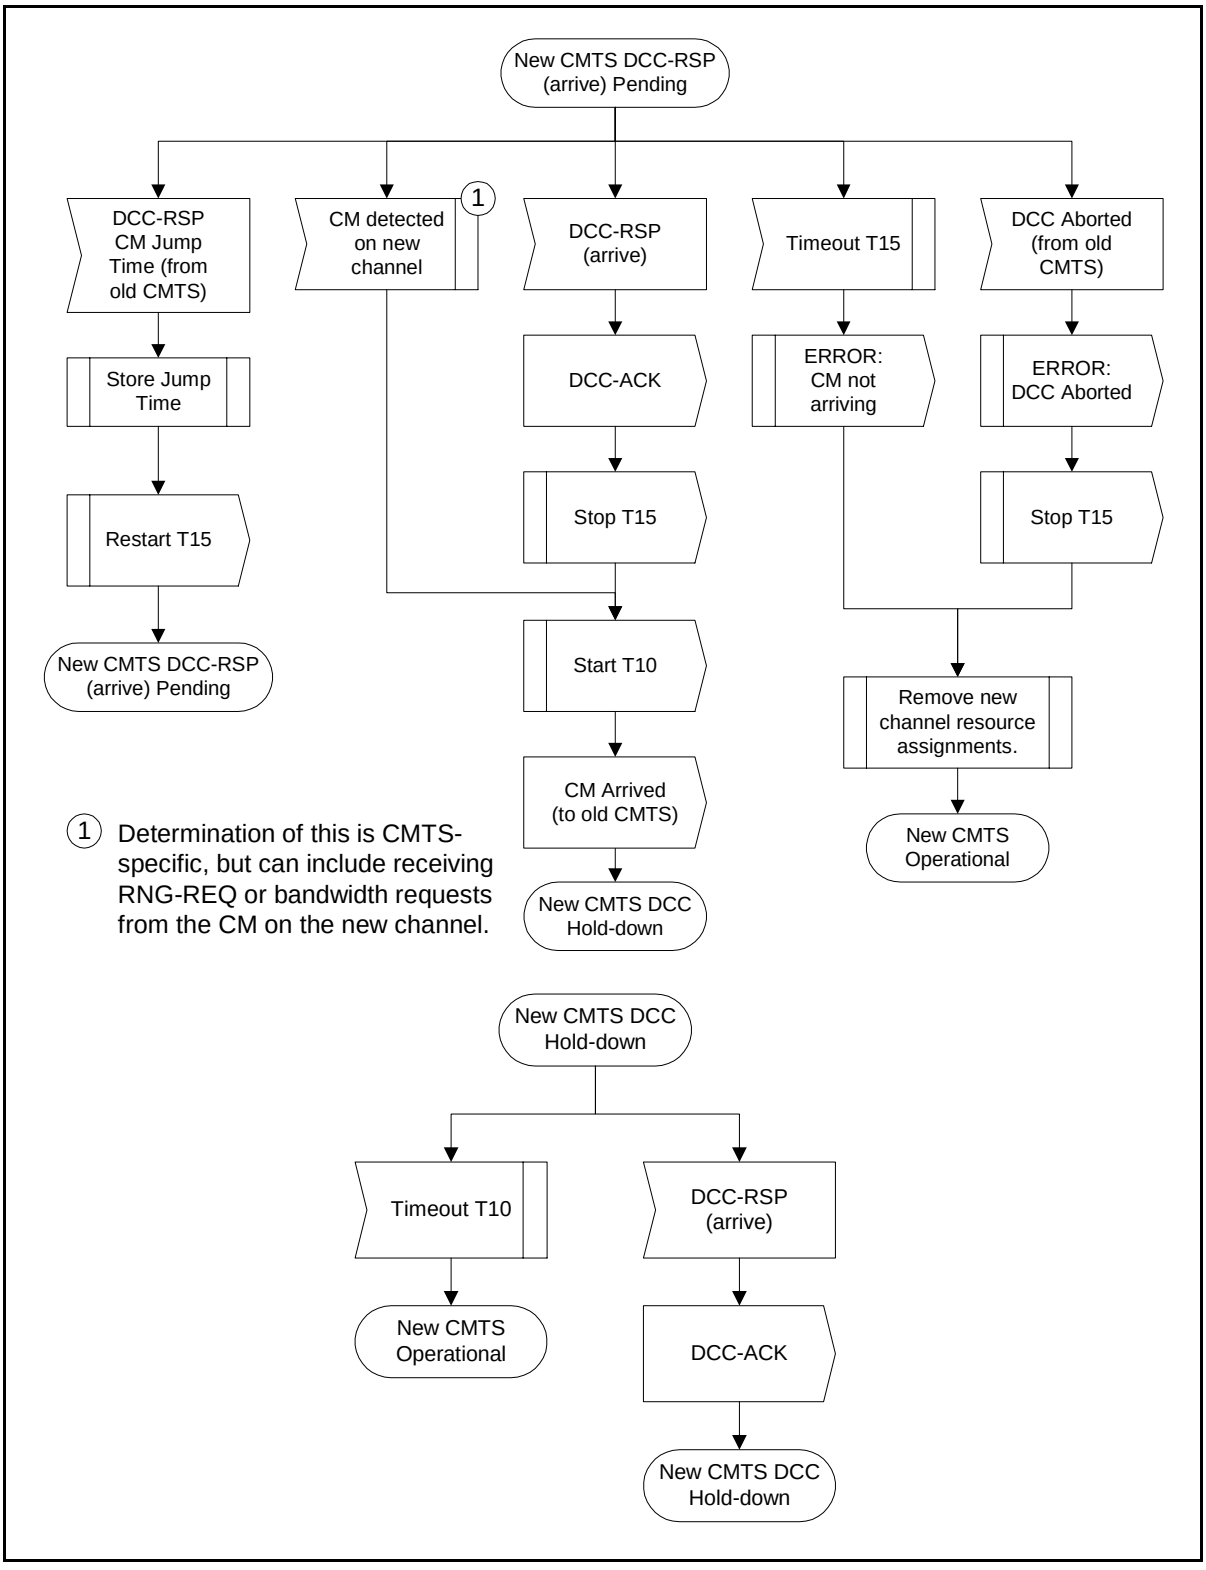


Figure 9-61. Dynamically Changing Channels: CMTS View Part 4

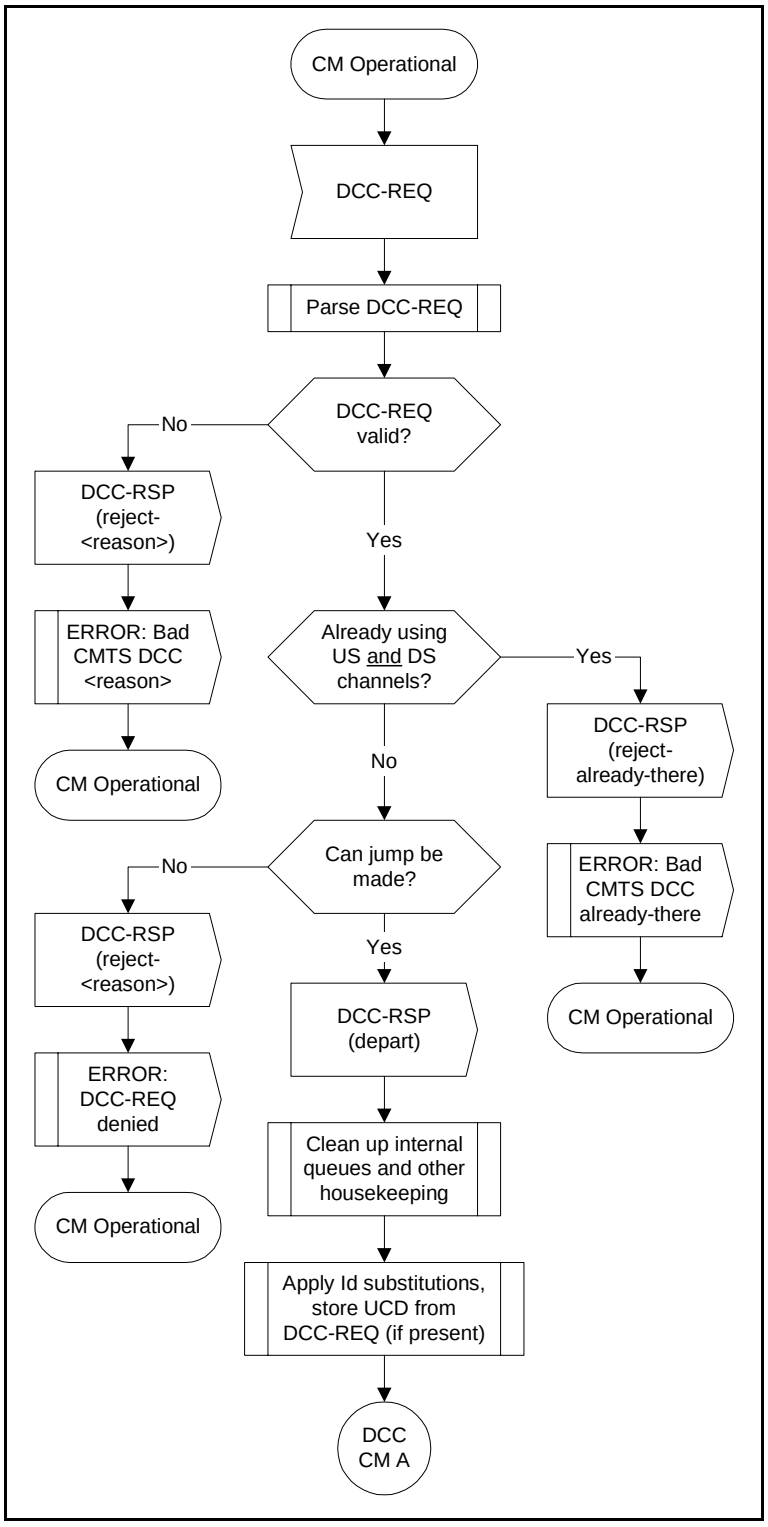


Figure 9-62. Dynamically Changing Channels: CM View Part 1[2]

2. The state "Obtain Upstream Parameters" links to the state machine in Figure 9-1.

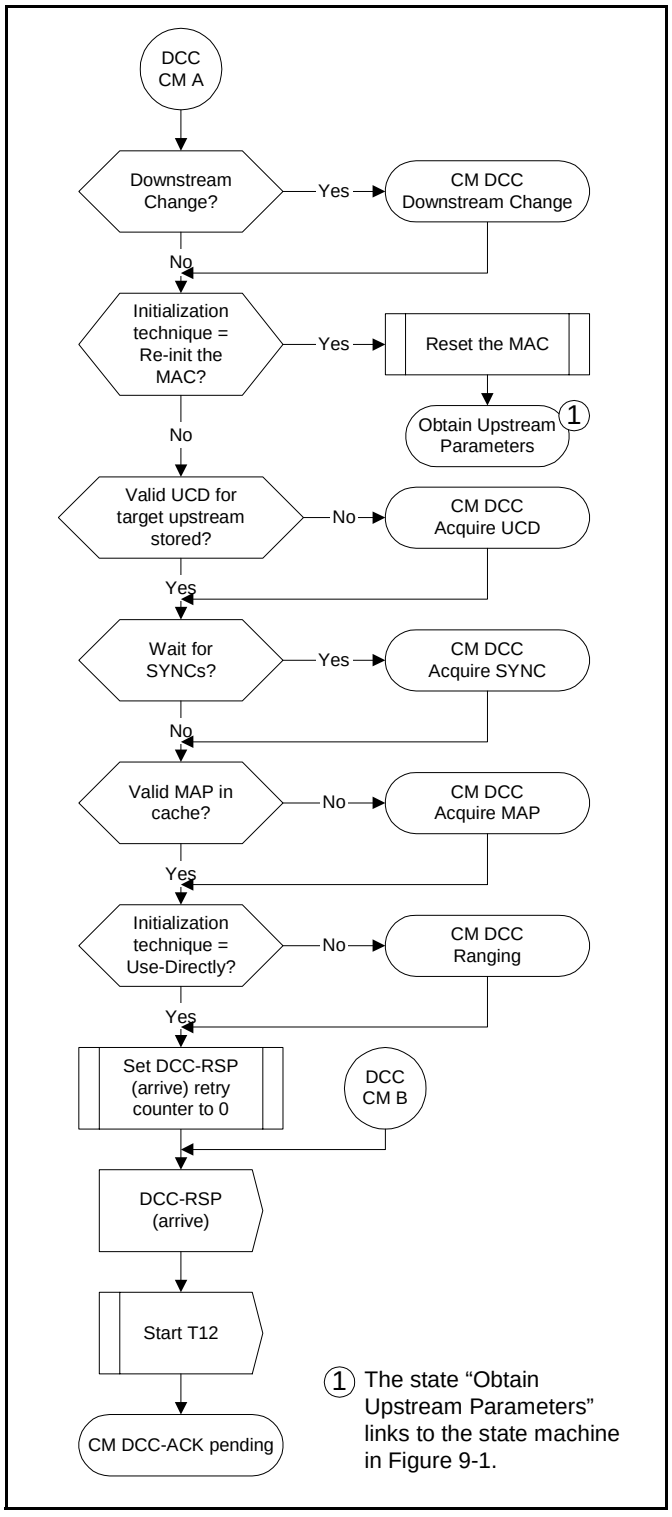


Figure 9-63. Dynamically Changing Channels: CM View Part 2

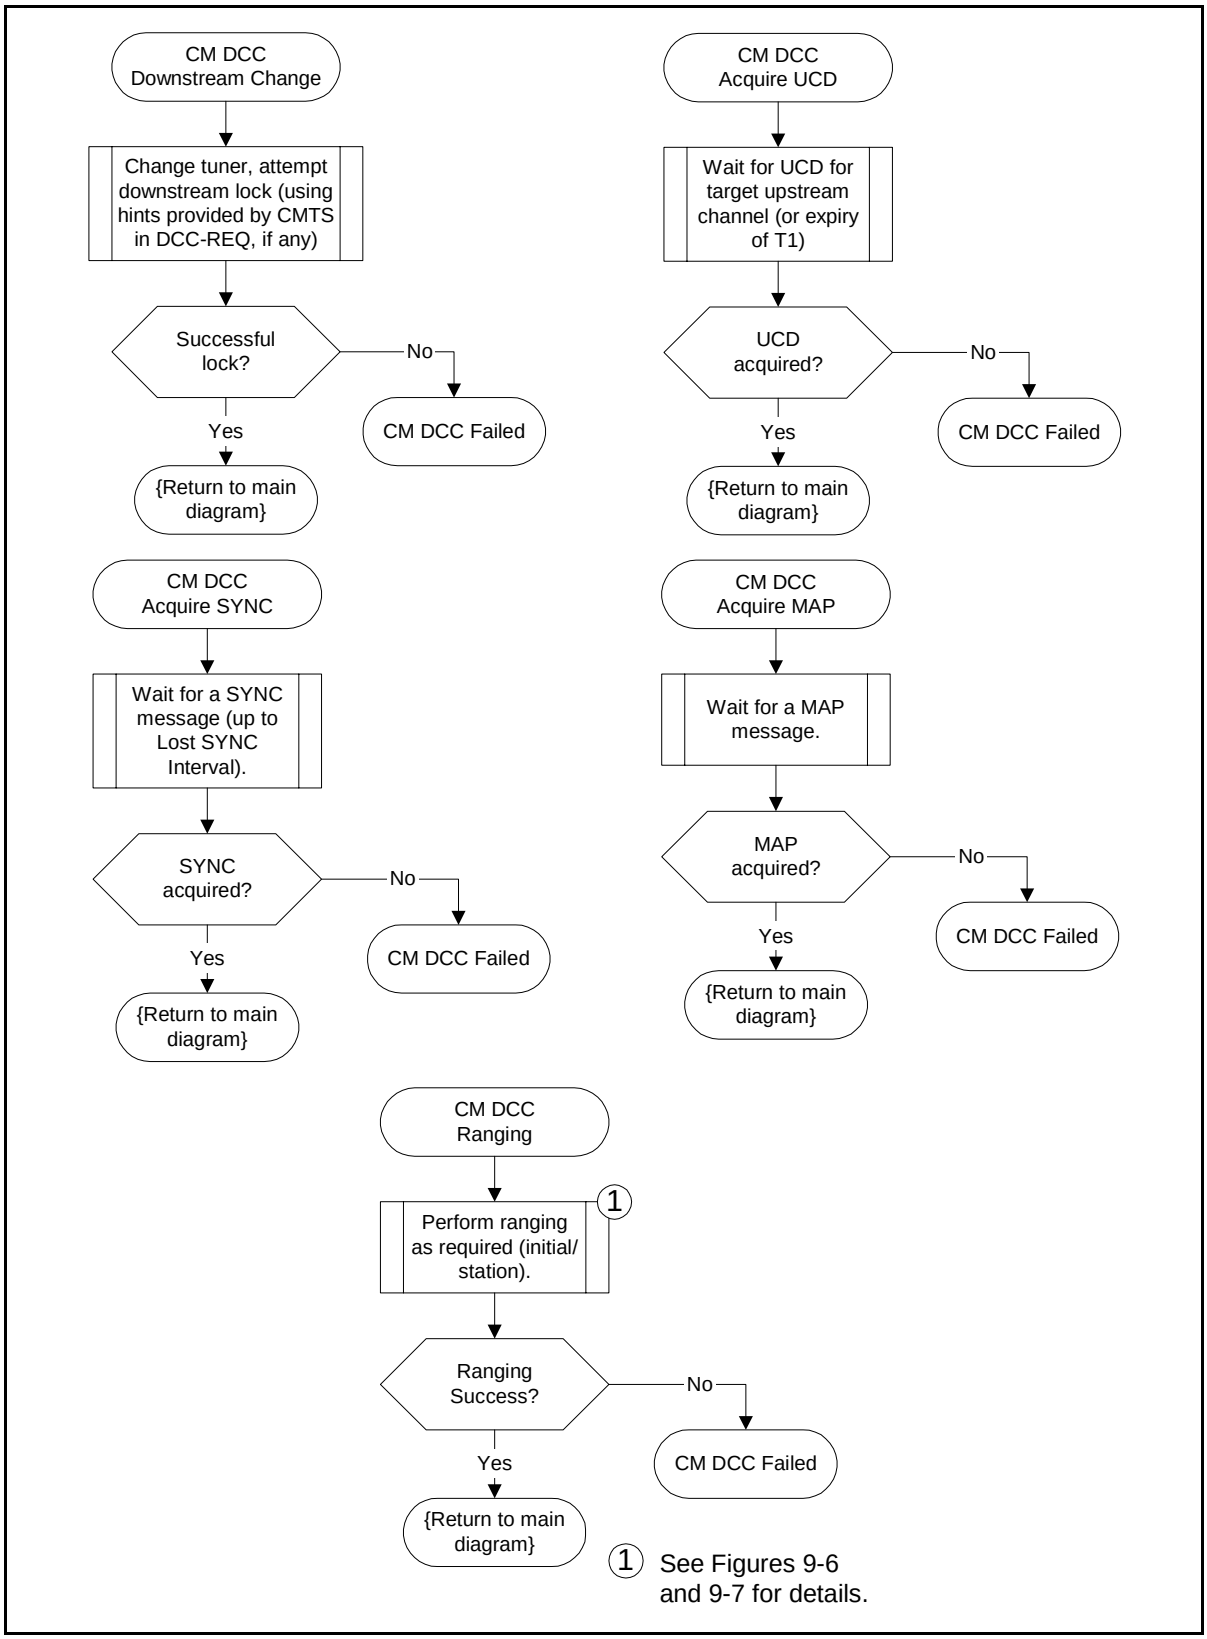


Figure 9-64. Dynamically Changing Channels: CM View Part 3

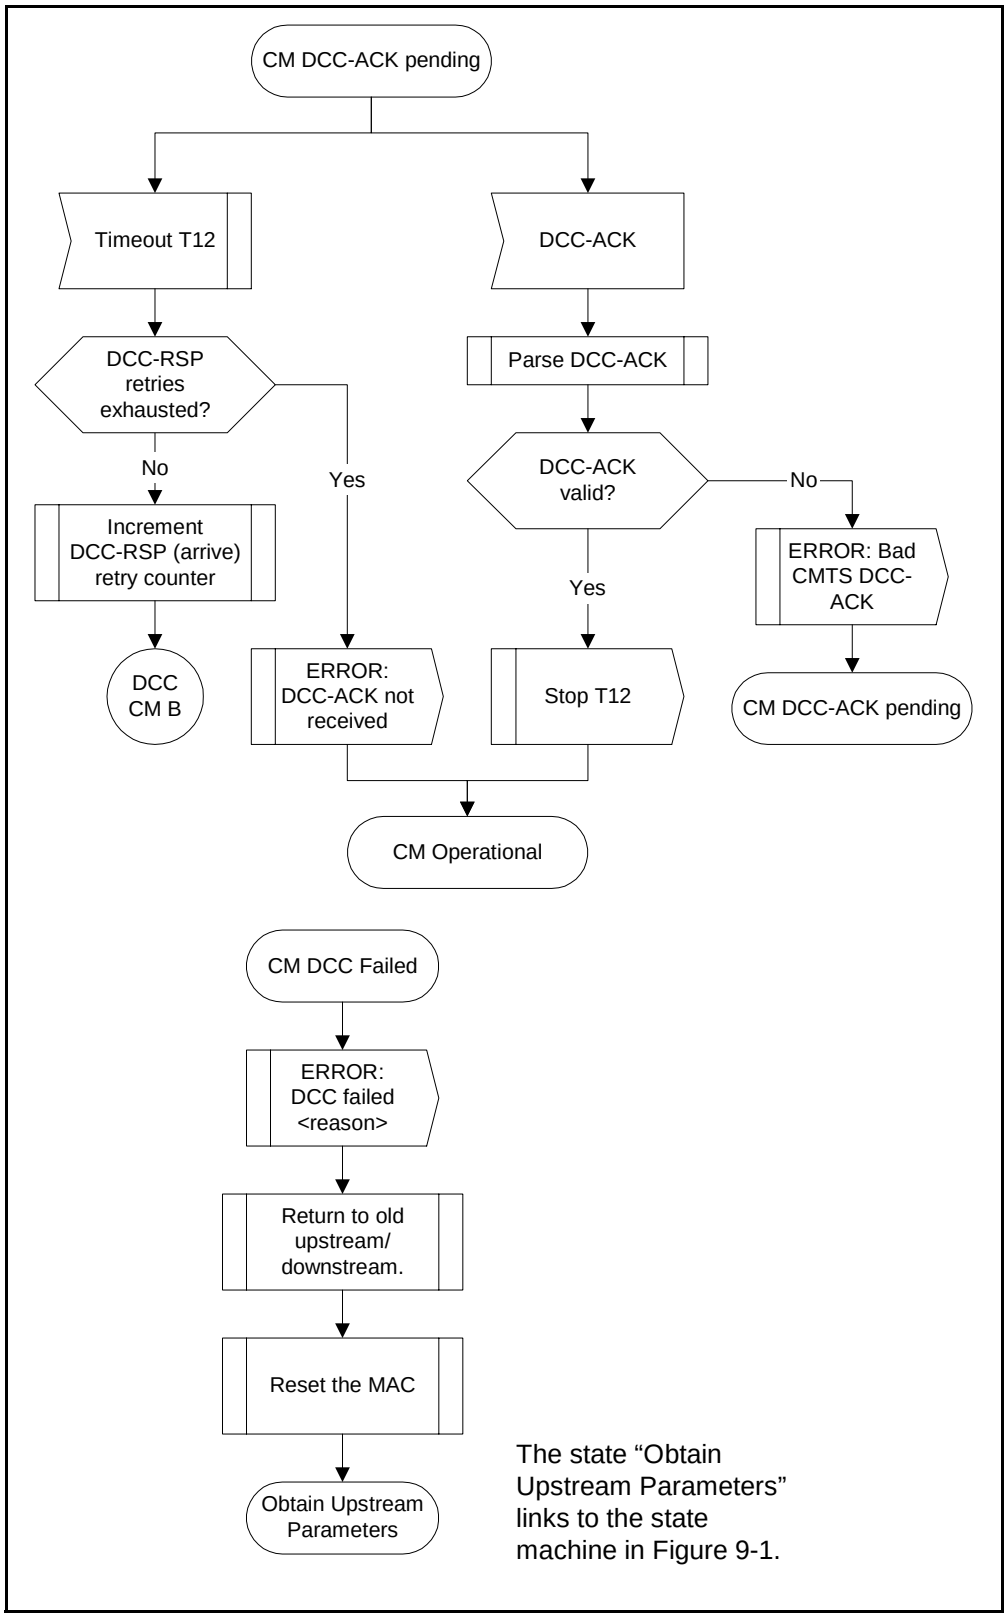


Figure 9-65. Dynamically Changing Channels: CM View Part 4

9.4.5.5.2 Example Timing

9.4.5.5.2.1 Upstream and Downstream Change – Use Channel Directly: CMTS Supplies All TLV Hints

In this example, the current CMTS sends a DCC-REQ message requesting that the CM switch both upstream and downstream channels. The DCC-REQ message includes the UCD substitution TLV, the SYNC substitution TLV, the downstream parameter TLVs, and the initialization technique TLV of 4 (use channel directly). The CM does not include the CM jump time TLV in the DCC-RSP.

The destination CMTS has the following local parameters:

UCD interval – 1 sec.
SYNC interval – 10 msec.
Unicast ranging interval – 1 sec.

The destination CMTS calculates the T15 timer value. The definition of the formula used in determining T15 is shown below. The variables used in calculating T15 are explained in the table below.

T15 = CmJumpTime + CmRxTargetUcd + CmRxDsSync + CmtsRxRngReq
T15 = 1.3 sec. + 2 sec. + 20 msec. + (2.02 sec.) = 5.34 sec.

Table 9-1. Use Channel Directly: CMTS Supplies All TLV Hints

Variable	Value	Explanation
CmJumpTime	1.3 sec.	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1.3 seconds.
CmRxTargetUcd	2 sec.	Although UCD substitution settings are specified in the DCC-REQ, the CMTS does not know that the CM implements this TLV.
CmRxDsSync	20 msec.	Although SYNC substitution settings are specified in the DCC-REQ, the CMTS does not know that the CM implements this TLV.
CmtsRxRngReq	2.02 sec = 2 * (1 sec.) + 20 msec.	Two times the CMTS time period between unicast ranging opportunities plus 20 - 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.

The CM synchronizes to the downstream parameters on the new channel, applies the UCD supplied in the DCC-REQ, collects MAP messages on the new channel, and resumes normal data transmission on the destination channels. This occurs within the recommended performance of 1 second.

9.4.5.5.2.2 Upstream and Downstream Change - Station maintenance: CMTS Supplies No TLV Hints

In this example, the current CMTS sends a DCC-REQ message requesting that the CM switch both upstream and downstream channels. The DCC-REQ message includes the initialization technique TLV of 2 (perform station maintenance). The CM does not include the CM jump time TLV in the DCC-RSP.

The destination CMTS has the following local parameters:

UCD interval – 1 sec.
SYNC interval – 10 msec.
Unicast ranging interval – 1 sec.

The destination CMTS starts scheduling the CM immediately after it receives the DCC-RSP (depart). The destination CMTS calculates the T15 timer value. The definition of the formula used in determining T15 is shown below. The variables used in calculating T15 are explained in the table below.

T15 = CmJumpTime + CmRxTargetUcd + CmRxDsSync + CmtsRxRngReq
T15 = 1.3 sec. + 2 sec. + 20 msec. + (2.02 sec.) = 5.34 sec.

Table 9-2. Station maintenance: CMTS Supplies No TLV Hints

Variable	Value	Explanation
CmJumpTime	1.3 sec.	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1.3 seconds.
CmRxTargetUcd	2 sec.	Two CMTS UCD timer periods.
CmRxDsSync	20 msec.	Two CMTS SYNC timer periods.
CmtsRxRngReq	2.02 sec = 2 * (1 sec.) + 20 msec.	Two times the CMTS time period between unicast ranging opportunities plus 20 - 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.

The CM should synchronize to the downstream parameters on the new channel, search for and apply a UCD message on the destination channel, wait for a downstream SYNC on the destination channel, collect MAP messages on the destination channel, perform station maintenance on the destination channel, and resume normal data transmission on the destination channels.

These events occur in less than two seconds; this is within the acceptable performance criteria. The DCC transaction occurred within the recommended four second sum of CM jump time, two UCD intervals, and two ranging intervals (0 + 2 sec. + 2 sec. = 4 sec.).

9.5 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 9-3 shows the recovery process that MUST be taken following the loss of a specific type of MAC message.

SP-OSSIv1.1 [DOCSIS5] Appendix F contains a list of error codes with more useful information as to the failure of the PHY and MAC layers. Refer to Section 6.2.8 for additional information.

Table 9-3. 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 Lost SYNC interval (see Appendix Appendix B) before it has lost synchronization with the network. A CM that has lost synchronization MUST NOT use the upstream and MUST try to re-establish synchronization.
UCD	During CM initialization the CM MUST receive a usable UCD[1] before transmitting on the upstream. When in the "Obtain Upstream Parameters" state of CM initialization process, if the CM doesn't receive a usable 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 Appendix Appendix 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 Appendix Appendix 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 Appendix Appendix B). Failure to receive a valid response after the requisite number of attempts MUST cause the CMTS to consider the CM as unreachable.

Notes to Table 9-1:

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.

9.5.1 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.

This page intentionally left blank.

10 Supporting Future New Cable Modem Capabilities

10.1 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 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 mechanism used for download MUST be TFTP file transfer. The mechanism by which transfers are secured and authenticated is in [DOCSIS8]. The transfer MUST be initiated in one of two ways:

- An SNMP manager requests the CM to upgrade.
- 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 parameter. If present, the CM MUST attempt to download the specified file from this server. If not present, the CM MUST attempt to download the specified file from the configuration file server.

The CM MUST verify that the downloaded image is appropriate for itself. If the image is appropriate, the CM MUST write the new software image to non-volatile storage. Once the file transfer is completed successfully, the CM MUST restart itself with the new code image.

If the CM is unable to complete the file transfer for any reason, it MUST remain capable of accepting new software downloads (without operator or user interaction), even if power or connectivity is interrupted between attempts. The CM MUST log the failure and MAY report it asynchronously to the network manager.

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, then 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, then it is critical that it MUST inter-work with the previous version in order to allow a gradual transition of units on the network.

This page intentionally left blank.

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 address range

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

A.2.1 All CMs and No CM Service IDs

These Service IDs are used in MAPs for special purposes or to indicate that any CM can respond in the corresponding interval.

0x0000	Addressed to no CM. Typically used when changing upstream burst parameters so that CMs have time to adjust their modulators before the new upstream settings are in effect. This is also the "Initialization SID" used by the CM during initial ranging."
0x3FFF	Addressed to all CMs. Typically used for broadcast Request intervals or Initial Maintenance intervals.

A.2.2 Well-Known 'Multicast' Service IDs

These Service IDs are only used for Request/Data IE's. They indicate that any CM can respond in a given interval, but that it must limit the size of its transmission to a particular number of minislots (as indicated by the particular multicast SID assigned to the interval).

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.2.3 Priority Request Service IDs

These Service IDs (0x3Exx) are reserved for Request IEs (refer to Appendix C.2.2.5.1).

If 0x01 bit is set, priority zero can request
If 0x02 bit is set, priority one can request
If 0x04 bit is set, priority two can request
If 0x08 bit is set, priority three can request
If 0x10 bit is set, priority four can request
If 0x20 bit is set, priority five can request
If 0x40 bit is set, priority six can request
If 0x80 bit is set, priority seven can request

Bits can be combined as desired by the CMTS upstream scheduler for any Request IUCs.

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)			200 msec
CMTS	UCD Interval	Time between transmission of UCD messages (ref. 6.3.3)			2 sec
CMTS	Max MAP Pending	The number of mini-slots that a CMTS is allowed to map into the future (ref. 6.3.4)			4096 mini-slot times
CMTS	Ranging Interval	Time between transmission of broadcast Ranging requests (ref. 7.3.3)			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. 9.2.4)	16		
CM, CMTS	Invited Ranging Retries	Number of Retries on inviting Ranging Requests (ref. 9.2.4)	16		
CM	Request Retries	Number of retries on bandwidth allocation requests	16		
CM CMTS	Registration Request/ Response Retries	Number of retries on registration requests/responses	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. 7.1.1)	200 µs		
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.	30 sec		
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, then 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 CMTS	T6	Wait for REG-RSP and REG-ACK			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		

System	Name	Time Reference	Minimum Value	Default Value	Maximum Value
CM CMTS	Timebase Tick	System timing unit	6.25 µsec		
CM CMTS	DSx Request Retries	Number of Timeout Retries on DSA/DSC/ DSD Requests	3		
CM CMTS	DSx Response Retries	Number of Timeout Retries on DSA/DSC/ DSD Responses	3		
CM CMTS	T7	Wait for DSA/DSC/DSD Response timeout			1 sec
CM CMTS	T8	Wait for DSA/DSC Acknowledge timeout			300 msec
CM	TFTP Backoff Start	Initial value for TFTP backoff	1sec		
CM	TFTP Backoff End	Last value for TFTP backoff	16 sec		
CM	TFTP Request Retries	Number of retries on TFTP request	16		
CM	TFTP Download Retries	Number of retries on entire TFTP downloads	3		
CM	TFTP Wait	The wait between TFTP retry sequences	10 min		
CM	ToD Retries	Number of Retries per ToD Retry Period	3		
CM	ToD Retry Period	Time period for ToD retries	5 min		
CMTS	T9	Registration Timeout, the time allowed between the CMTS sending a RNG-RSP (success) to a CM, and receiving a REG-REQ from that same CM.	15 min	15 min	
CM CMTS	T10	Wait for Transaction End timeout			3 sec
CMTS	T11	Wait for a DCC Response on the old channel			300 ms
CM	T12	Wait for a DCC Acknowledge			300 ms
CMTS	T13	Maximum holding time for QoS resources for DCC on the old channel			1 sec
CM	T14	Minimum time after a DSx reject-temp-DCC and the next retry of DSx command	2 sec		
CMTS	T15	Maximum holding time for QoS resources for DCC on the new channel	2 sec		35 sec
CMTS	DCC-REQ Retries	Number of retries on Dynamic Channel Change Request	3		
CM	DCC-RSP Retries	Number of retries on Dynamic Channel Change Response	3		
CM	Lost DCI-REQ interval	Time from sending DCI-REQ and not receiving a DCI-RSP			2 sec
CM	DCI-REQ retry	Number of retries of DCI-REQ before rebooting			16
CM	DCI Backoff start	Initial value for DCI backoff	1 sec		
CM	DCI Backoff end	Last value for DCI backoff	16 sec		

Appendix C Common Radio Frequency Interface Encodings

C.1 Encodings for Configuration and MAC-Layer Messaging

The following type/length/value encodings MUST be used in both the configuration file (see Appendix D), in CM registration requests and in Dynamic Service Messages. 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.1.1 Configuration File and Registration Settings

These settings are found in the configuration file and, if present, MUST be forwarded by the CM to the CMTS in its Registration Request.

C.1.1.1 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	Value
1	4	Rx Frequency

Valid Range:
The receive frequency MUST be a multiple of 62500 Hz.

C.1.1.2 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.1.1.3 Network Access Control Object

If the value field is a 1, CPE attached to this CM are allowed access to the network, based on CM provisioning. If the value of this field is a 0, the CM MUST NOT forward traffic from attached CPE to the RF MAC network, but MUST continue to accept and generate traffic from the CM itself. The value of this field does not affect CMTS service flow operation and does not affect CMTS data forwarding operation.

Type	Length	On / Off
3	1	1 or 0

Note: The intent of “NACO = 0” is that the CM does not forward traffic from any attached CPE onto the cable network. (A CPE is any client device attached to that CM, regardless of how that attachment is implemented.) However, with “NACO = 0”, management traffic to the CM is not restricted. Specifically, with NACO off, the CM remains manageable, including sending/receiving management traffic such as (but not limited to):

- ARP: allow the modem to resolve IP addresses, so it can respond to queries or send traps.
- DHCP: allow the modem to renew its IP address lease.
- ICMP: enable network troubleshooting for tools such as “ping” and “traceroute.”
- ToD: allow the modem to continue to synchronize its clock after boot.
- TFTP: allow the modem to download either a new configuration file or a new software image.
- SYSLOG: allow the modem to report network events.
- SNMP: allow management activity

In DOCSIS v1.1, with NACO off, the primary upstream and primary downstream service flows of the CM remain operational only for management traffic to and from the CM. With respect to DOCSIS v1.1 provisioning, a CMTS should ignore the NACO value and allocate any service flows that have been authorized by the provisioning server.

C.1.1.4 DOCSIS 1.0 Class of Service Configuration Setting

This field defines the parameters associated with a DOCSIS 1.0 class of service. Any CM registering with a DOCSIS 1.0 Class of Service Configuration Setting MUST be treated as a DOCSIS 1.0 CM. Refer to Section 6.3.8.

This field defines the parameters associated with a class of service. It is somewhat complex in that 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.1.1.4.1 Class ID

The value of the field specifies the identifier for the class of service to which the encapsulated string applies.

Type	Length	Value
4.1	1	

Valid Range

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

C.1.1.4.2 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
4.2	4	

Note: This is a limit, not a guarantee that this rate is available.

C.1.1.4.3 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
4.3	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.1.1.4.4 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
4.4	1	

Valid Range
0 -> 7

C.1.1.4.5 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
4.5	4	

C.1.1.4.6 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
4.6	2	

C.1.1.4.7 Class-of-Service Privacy Enable

This configuration setting enables/disables Baseline Privacy on a provisioned CoS. See [SCTE2].

Type	Length	Enable / Disable
4.7 (= CoS_BP_ENABLE)	1	1 or 0

Table C-1. Sample DOCSIS 1.0 Class of Service Encoding

Type	Length	Value (sub)type	Length	Value	
4	28				class of service configuration setting
		1	1	1	service class
		2	4	10,000,000	max. downstream rate of 10 Mb/sec
		3	4	300,000	max. upstream rate of 300 kbps
		4	1	5	return path priority of 5
		5	4	64,000	min guaranteed 64 kb/sec
		6	2	1518	max. Tx burst of 1518 bytes
4	28				class of service configuration setting
		1	1	2	service class 2
		2	4	5,000,000	max. forward rate of 5 Mb/sec
		3	4	300,000	max. return rate of 300 Mb/sec
		4	1	3	return path priority of 3
		5	4	32,000	min guaranteed 32 kb/sec
		6	2	1,518	max. Tx burst of 1518 bytes

C.1.1.5 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.1.1.6 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.1.1.7 Maximum Number of CPEs

The maximum number of CPEs that can be granted access 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. The maximum number of CPE's MUST be enforced by the CM.

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.

Type	Length	Value
18	1	

The CM MUST interpret this value as an unsigned integer. The non-existence of this option, or the value 0, MUST be interpreted as 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 use a lower value.

C.1.1.8 TFTP Server Timestamp

The sending time of the configuration file in seconds. The definition of time is as in [RFC-868]

Type	Length	Value
19	4	Number of seconds since 00:00 1 Jan 1900

Note: The purpose of this parameter is to prevent replay attacks with old configuration files.

C.1.1.9 TFTP Server Provisioned Modem Address

The IP Address of the modem requesting the configuration file.

Type	Length	Value
20	4	IP Address

Note: The purpose of this parameter is to prevent IP spoofing during registration.

C.1.1.10 Upstream Packet Classification Configuration Setting

This field defines the parameters associated with one entry in an upstream traffic classification list. Refer to Appendix C.2.1.1.

Type	Length	Value
22	n	

C.1.1.11 Downstream Packet Classification Configuration Setting

This field defines the parameters associated with one Classifier in an downstream traffic classification list. Refer to Appendix C.2.1.2.

Type	Length	Value
23	n	

C.1.1.12 Upstream Service Flow Encodings

This field defines the parameters associated with upstream scheduling for one Service Flow. Refer to Appendix C.2.2.1.

Type	Length	Value
24	n	

C.1.1.13 Downstream Service Flow Encodings

This field defines the parameters associated with downstream scheduling for one Service Flow. Refer to Appendix C.2.2.2.

Type	Length	Value
25	n	

C.1.1.14 Payload Header Suppression

This field defines the parameters associated with Payload Header Suppression.

Type	Length	Value
26	n	

C.1.1.15 Maximum Number of Classifiers

This is the maximum number of Classifiers associated with admitted or active upstream Service Flows that the CM is allowed to have. Both active and inactive Classifiers are included in the count.

This is useful when using deferred activation of provisioned resources. The number of provisioned Service Flows may be high and each Service Flow might support multiple Classifiers. Provisioning represents the set of Service Flows the CM can choose between. The CMTS can control the QoS resources committed to the CM by limiting the number of Service Flows that are admitted. However, it may still be desirable to limit the number of Classifiers associated with the committed QoS resources. This parameter provides that limit.

Type	Length	Value
28	2	Maximum number of active and inactive Classifiers associated with admitted or active upstream Service Flows

The default value MUST be 0 — no limit.

C.1.1.16 Privacy Enable

This configuration setting enables/disables Baseline Privacy [DOCSIS8] on the Primary Service Flow and all other Service Flows for this CM. If a DOCSIS 1.1 CM receives this setting in a configuration file, the CM is required to forward this setting as part of the registration request (REG-REQ) as specified in Section 6.3.7 regardless of whether the configuration file is DOCSIS 1.1-style or not while this setting is usually contained only in a DOCSIS 1.1-style configuration file with DOCSIS 1.1 Service Flow TLVs.

Type	Length	Value
29	1	0 — Disable 1 — Enable

The default value of this parameter MUST be 1 — privacy enabled.

C.1.1.17 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 (Appendix C.1.3.2) 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, then the TLV MUST be discarded.

This configuration setting MAY appear multiple times. The same Vendor ID MAY appear multiple times. This configuration setting MAY be nested inside a Packet Classification Configuration Setting, a Service Flow Configuration Setting, or a Service Flow Response. 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.1.1.18 Subscriber Management TLVs

The information in these TLVs is not used by the CM; rather, the information is used by the CMTS to populate the Subscriber Management MIB for this CM.

If present in the configuration file, the CM MUST include these TLVs in the subsequent REG-REQ to be used by the CMTS to populate the Subscriber Management MIB for this CM. If present in the configuration file, the CM MUST include these TLVs in the CMTS MIC.

C.1.1.18.1 Subscriber Management Control

This three byte field provides control information to the CMTS for the Subscriber Management MIB. The first two bytes represent the number of IP addresses permitted behind the CM. The third byte is used for control fields.

Type	Length	Value
35	3	byte 1,2 docsSubMgtCpeControlMaxCpeIP (low order 10 bits) byte 3, bit 0: docsSubMgtCpeControlActive byte 3, bit 1: docsSubMgtCpeControlLearnable byte 3, bits #2-7: reserved, must be set to zero

C.1.1.18.2 Subscriber Management CPE IP Table

This field lists the IP Addresses used to populate docsSubMgtCpeIpTable in the Subscriber Management MIB at the CMTS.

Type	Length	Value
36	N (multiple of 4)	Ipa1, Ipa2, Ipa3, Ipa4

C.1.1.18.3 Subscriber Management Filter Groups

The Subscriber Management MIB allows filter groups to be assigned to a CM and CPE attached to that CM. These include two CM filter groups, upstream and downstream, and two CPE filter groups, upstream and downstream. These four filter groups are encoded in the configuration file in a single TLV as follows:

Type	Length	Value
37	8	bytes 1,2: docsSubMgtSubFilterDownstream group bytes 3,4: docsSubMgtSubFilterUpstream group bytes 5,6: docsSubMgtCmFilterDownstream group bytes 7,8: docsSubMgtCmFilterUpstream group

C.1.2 Configuration-File-Specific Settings

These settings are found in only the configuration file. They MUST NOT be forwarded to the CMTS in the Registration Request.

C.1.2.1 End-of-Data Marker

This is a special marker for end of data.

It has no length or value fields.

Type
255

C.1.2.2 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.1.2.3 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 D.2.2. See Section 10.1.

Type	Length	Value
9	n	filename

C.1.2.4 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 [ISO/IEC 8825-1] 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.1.2.5 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.1.2.6 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

C.1.2.7 Software Upgrade TFTP Server

The IP address of the TFTP server, on which the software upgrade file for the CM resides. See Section 10.1 and Appendix C.1.2.3

Type	Length	Value
21	4	ip1,ip2,ip3,ip4

C.1.2.8 SnmpV3 Kickstart Value

Compliant CMs MUST understand the following TLV and its sub-elements and be able to kickstart SNMPv3 access to the CM regardless of whether the CMs are operating in 1.0 mode or 1.1 mode

Type	Length	Value
34	n	Composite

Up to 5 of these objects may be included in the configuration file. Each results in an additional row being added to the usmDHKickstartTable and the usmUserTable and results in an agent public number being generated for those rows.

C.1.2.8.1 SnmpV3 Kickstart Security Name

Type	Length	Value
34.1	2-16	UTF8 Encoded security name

For the ASCII character set, the UTF8 and the ASCI I encodings are identical. Normally, this will be specified as one of the DOCSIS built-in USM users, *e.g.*, "docsisManager," "docsisOperator," "docsisMonitor," "docsisUser." The security name is NOT zero terminated. This is reported in the usmDHKickStartTable as usmDHKickStartSecurityName and in the usmUserTable as usmUserName and usmUserSecurityName.

C.1.2.8.2 SnmpV3 Kickstart Manager Public Number

Type	Length	Value
34.2	n	Manager's Diffie-Helman public number expressed as an octet string.

This number is the Diffie-Helman public number derived from a privately (by the manager or operator) generated random number and transformed according to [RFC-2786]. This is reported in the usmDHKickStartTable as usmKickstartMgrPublic. When combined with the object reported in the same row as usmKickstartMyPublicit can be used to derive the keys in the related row in the usmUserTable.

C.1.2.9 Manufacturer Code Verification Certificate

The Manufacturer's Code Verification Certificate (M-CVC) for Secure Software Downloading specified by the Appendix D of [DOCSIS8]. The CM config file MUST contain this M-CVC and/or C-CVC defined in C.1.2.10 in order to allow the 1.1 compliant CM to download the code file from TFTP server regardless the CM is provisioned to run with BPI, BPI+, or none of them. See [DOCSIS8] Appendix D for detail.

Type	Length	Value
32	n	Manufacturer CVC (DER-encoded ASN.1)

If the length of the M-CVC exceeds 254 bytes, the M-CVC MUST be fragmented into two or more successive Type 32 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the M-CVC by concatenating the contents (Value of the TLV) of successive Type 32 elements in the order in which they appear in the config file. For example, the first byte following the length field of the second Type 32 element is treated as if it immediately follows the last byte of the first Type 32 element.

C.1.2.10 Co-signer Code Verification Certificate

The Co-signer's Code Verification Certificate (C-CVC) for Secure Software Downloading specified by the Appendix D of [DOCSIS8]. The CM config file MUST contain this C-CVC and/or M-CVC defined in C.1.2.9 in order to allow the 1.1 compliant CM to download the code file from TFTP server regardless the CM is provisioned to run with BPI, BPI+, or none of them. See [DOCSIS8] Appendix D for detail.

Type	Length	Value
33	n	Co-signer CVC (DER-encoded ASN.1)

If the length of the C-CVC exceeds 254 bytes, the C-CVC MUST be fragmented into two or more successive Type 33 elements. Each fragment, except the last, MUST be 254 bytes in length. The CM reconstructs the C-CVC by concatenating the contents (Value of the TLV) of successive Type 33 elements in the order in which they appear in the config file. For example, the first byte following the length field of the second Type 33 element is treated as if it immediately follows the last byte of the first Type 33 element.

C.1.2.11 SNMPv3 Notification Receiver

This TLV specifies a Network Management Station that will receive notifications from the modem when it is in Coexistence mode.

Type	Length	Value
38	n	Composite

C.1.2.11.1 SNMPv3 Notification Receiver IP Address

This sub-TLV specifies the IP address of the notification receiver.

Type	Length	Value
38.1	4	ip1, ip2, ip3, ip4

If TLV 38.1 is not present, the CM MUST consider this a configuration failure, and the CM MUST NOT proceed with CM registration.

C.1.2.11.2 SNMPv3 Notification Receiver UDP Port Number

This sub-TLV specifies the Port number on the notification receiver to receive the notifications.

Type	Length	Value
38.2	2	UDP port number

If not present, the default value 162 is used.

C.1.2.11.3 SNMPv3 Notification Receiver Trap Type

This sub-TLV specifies the type of trap to send.

Type	Length	Value
38.3	2	1 : SNMP v1 trap in an SNMP v1 packet 2 : SNMP v2c trap in an SNMP v2c packet 3 : SNMP inform in an SNMP v2c packet 4 : SNMP v2c trap in an SNMP v3 packet 5 : SNMP inform in an SNMP v3 packet

If TLV 38.3 is not present, the CM MUST consider this a configuration failure, and the CM MUST NOT proceed with CM registration.

C.1.2.11.4 SNMPv3 Notification Receiver Timeout

This sub-TLV specifies the round trip timeout used to wait before sending a retry of an inform notification if sender does not get an acknowledgement from the receiver.

Type	Length	Value
38.4	2	time in milliseconds

If not present, the default value of 15000 milliseconds is used. This corresponds to the default value of 1500 hundredths of a second defined for the snmpTargetAddrTimeout MIB object (see Appendix P of [DOCSIS5] and [RFC-2573]).

C.1.2.11.5 SNMPv3 Notification Receiver Retries

Defines the number times to retry an Inform after the first Inform transmission.

Type	Length	Value
38.5	2	number of retries

If not present, the default value of 3 retries is used.

SNMPv3 Notification Receiver Retries must be in the range of 0 to 255.

C.1.2.11.6 Notification Receiver Filtering Parameters

This sub-TLV specifies the OID of the snmpTrapOID value that is the root of the MIB subtree that defines all of the notifications to be sent to the Notification Receiver.

Type	Length	Value
38.6	n	Object Identifier ASN.1

The encoding of this TLV value field starts with the ASN.1 Universal type 6 (Object Identifier) byte, then the ASN.1 length field, then the ASN.1 encoded object identifier components. If this Sub-TLV is not present, the notification receiver will receive all notifications generated by the SNMP agent.

C.1.2.11.7 Notification Receiver Security Name

This sub-TLV specifies v3 Security Name to use when sending a SNMP V3 Notification.

Type	Length	Value
38.7	2-16	UTF8 Encoded security name

When Trap of Type value field is set to 1, 2, or 3, this value field SHOULD not be interpreted (has no meaning) and Informs messages will be sent with community string “public”. In the case of Trap of Type 4 or 5, two situations happen:

- If this TLV is not supplied, the V3 Notification will be sent in the noAuthNoPriv security level using the security name “@config”.
- If TLV-38 is supplied in configuration file, the value field MUST be the Security Name specified in a TLV Type 34 as part of the DH Kickstart procedure. The notifications will be sent using the Authentication and Privacy Keys calculated by the modem during the DH Kickstart procedure.

For detailed implementation refer to Section 3.6, Config File Element - docsisV3NotificationReceiver, of [DOCSIS5].

C.1.3 Registration-Request/Response-Specific Encodings

These encodings are not found in the configuration file, but are included in the Registration Request and option 60 of the DHCP request. Some encodings are also used in the Registration Response.

The CM MUST include all Modem Capabilities Encodings that are subject to negotiation with the CMTS in its Registration Request. Modem Capabilities Encodings that are not subject to negotiation with the CMTS are explicitly stated in the description of the particular modem capability. The CMTS MUST include Modem Capabilities in the Registration Response.

C.1.3.1 Modem Capabilities Encoding

The value field describes the capabilities of a particular modem, *i.e.*, implementation dependent limits on the particular features or number of features which the modem can support. It is composed from a number of encapsulated type/length/value fields. The encapsulated sub-types define the specific capabilities for the modem in question. Note that the sub-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.

All these capabilities are to be included in both the registration request and option 60 of the DHCP request unless the description of the capability explicitly prohibits this.

C.1.3.1.1 Concatenation Support

If the value field is a 1 the CM requests concatenation support from the CMTS.

Type	Length	On / Off
5.1	1	1 or 0

C.1.3.1.2 DOCSIS Version

DOCSIS version of this modem.

Type	Length	Value
5.2	1	0: DOCSIS v1.0 1: DOCSIS v1.1 2-255: Reserved

If this tuple is absent, the CMTS MUST assume DOCSIS v1.0 operation. The absence of this tuple or the value ‘DOCSIS 1.0’ does not necessarily mean the CM only supports DOCSIS 1.0 functionality — the CM MAY indicate it supports other individual capabilities with other Modem Capability Encodings. (Refer to Appendix G.3)

C.1.3.1.3 Fragmentation Support

If the value field is a 1 the CM requests fragmentation support from the CMTS.

Type	Length	Value
5.3	1	1 or 0

C.1.3.1.4 Payload Header Suppression Support

If the value field is a 1 the CM requests payload header suppression support from the CMTS.

Type	Length	Value
5.4	1	1 or 0

C.1.3.1.5 IGMP Support

If the value field is a 1 the CM supports DOCSIS 1.1-compliant IGMP.

Type	Length	Value
5.5	1	1 or 0

Note: This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a CMTS does receive this capability with in a registration request it MUST return the capability with the same value in the registration response.

C.1.3.1.6 Privacy Support

The value is the BPI support of the CM.

Type	Length	Value	
5.6	1	0	BPI Support
		1	BPI Plus Support
		2 - 255	Reserved

C.1.3.1.7 Downstream SAID Support

This field shows the number of Downstream SAIDs the modem can support.

Type	Length	Value
5.7	1	Number of Downstream SAIDs the CM can support.

If the number of SAIDs is 0 that means the Modem can support only 1 SAID.

C.1.3.1.8 Upstream SID Support

This field shows the number of Upstream SIDs the modem can support.

Type	Length	Value
5.8	1	Number of Upstream SIDs the CM can support.

If the number of SIDs is 0 that means the Modem can support only 1 SID.

C.1.3.1.9 Optional Filtering Support

This field shows the optional filtering support in the modem.

Type	Length	Value
5.9	1	Packet Filtering Support Array bit #0: 802.1P filtering bit #1: 802.1Q filtering bit #2-7: reserved MUST be set to zero

Note: This CM capability is not subject to negotiation with the CMTS. The CM MUST include this capability in the DHCP request, but MUST NOT include this capability in the registration request. If a CMTS does receive this capability with in a registration request it MUST return the capability with the same value in the registration response.

C.1.3.1.10 Transmit Equalizer Taps per Symbol

This field shows the maximal number of pre-equalizer taps per symbol supported by the CM.

Note: All CMs 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.1.3.1.11 Number of Transmit Equalizer Taps

This field shows the number of equalizer taps that are supported by the CM

Note: All CMs 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.1.3.1.12 DCC Support

The value is the DCC support of the CM.

Type	Length	Value
5.12	1	0 = DCC is not supported 1 = DCC is supported

C.1.3.2 Vendor ID Encoding

The value field contains the vendor identification specified by the three-byte vendor-specific Organization Unique Identifier of the CM MAC address.

The Vendor ID MUST be used in a Registration Request, but MUST NOT be used as a stand-alone configuration file element. It MAY be used as a sub-field of the Vendor Specific Information Field in a configuration file. When used as a sub-field of the Vendor Specific Information field, this identifies the Vendor ID of the CMs

which are intended to use this information. When the vendor ID is used in a Registration Request, then it is the Vendor ID of the CM sending the request.

Type	Length	Value
8	3	v1, v2, v3

C.1.3.3 Modem IP Address

For backwards compatibility with DOCSIS v 1.0. Replaced by 'TFTP Server Provisioned Modem Address'.

Type	Length	Value
12	4	IP Address

C.1.3.4 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, Confirmation Code

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
Confirmation Code	Refer to Appendix C.4.

C.1.4 Dynamic-Service-Message-Specific Encodings

These encodings are not found in the configuration file, nor in the Registration Request/Response signaling. They are only found in DSA-REQ, DSA-RSP, DSA-ACK, DSC-REQ, DSC-RSP, DSC-ACK, and DSD-REQ messages (Section 6.3.12 through 6.3.18).

C.1.4.1 HMAC-Digest

The HMAC-Digest setting is a keyed message digest. If privacy is enabled, the HMAC-Digest Attribute MUST be the final Attribute in the Dynamic Service message's Attribute list. The message digest is performed over the all of the Dynamic Service parameters (starting immediately after the MAC Management Message Header and up to, but not including the HMAC Digest setting), other than the HMAC-Digest, in the order in which they appear within the packet.

Inclusion of the keyed digest allows the receiver to authenticate the message. The HMAC-Digest algorithm, and the upstream and downstream key generation requirements are documented in [DOCSIS8].

This parameter contains a keyed hash used for message authentication. The HMAC algorithm is defined in [RFC-2104]. The HMAC algorithm is specified using a generic cryptographic hash algorithm. Baseline Privacy uses a particular version of HMAC that employs the Secure Hash Algorithm (SHA-1), defined in [SHA].

A summary of the HMAC-Digest Attribute format is shown below. The fields are transmitted from left to right.

Type	Length	Value
27	20	A 160-bit (20 octet) keyed SHA hash

C.1.4.2 Authorization Block

The Authorization Block contains an authorization "hint". The specifics of the contents of this "hint" are beyond the scope of this specification, but include [PKT-DQOS].

The Authorization Block MAY be present in CM-initiated DSA-REQ and DSC-REQ, and CMTS-initiated DSA-RSP and DSC-RSP messages. This parameter MUST NOT be present in CMTS-initiated DSA-REQ and DSC-REQ, nor CM-initiated DSA-RSP and DSC-RSP messages.

The Authorization Block information applies to the entire content of the message. Thus, only a single Authorization Block per message MAY be present. The Authorization Block, if present, MUST be passed to the Authorization Module in the CMTS. The Authorization Block information is only processed by the Authorization Module.

Type	Length	Value
30	n	Sequence of n octets

C.1.4.3 Key Sequence Number

The value shows the key sequence number of the BPI+ Authorization Key which is used to calculate the HMAC-Digest in case that the Privacy is enabled.

Type	Length	Value
31	1	Auth Key Sequence Number (0 - 15)

C.2 Quality-of-Service-Related Encodings

C.2.1 Packet Classification Encodings

The following type/length/value encodings MUST be used in both the configuration file, registration messages, and Dynamic Service messages to encode parameters for packet classification and scheduling. 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.

A classifier MUST contain at least one encoding from C.2.1.5 “IP Packet Classification Encodings”, C.2.1.6 “Ethernet LLC Packet Classification Encodings”, or C.2.1.7 “IEEE 802.1P/Q Packet Classification Encodings”.

The following configuration settings MUST be supported by all CMs which are compliant with this specification. All CMTSs MUST support classification of downstream packets based on IP header fields (Appendix C.2.1.5).

C.2.1.1 Upstream Packet Classification Encoding

This field defines the parameters associated with an upstream Classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream packet classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
22	n	

C.2.1.2 Downstream Packet Classification Encoding

This field defines the parameters associated with a downstream Classifier.

Note that the same subtype fields defined are valid for both the encapsulated upstream and downstream flow classification configuration setting string. These type fields are not valid in other encoding contexts.

Type	Length	Value
23	n	

C.2.1.3 General Packet Classifier Encodings

C.2.1.3.1 Classifier Reference

The value of the field specifies a reference for the Classifier. This value is unique per Dynamic Service message, configuration file, or Registration Request message.

Type	Length	Value
[22/23].1	1	1 - 255

C.2.1.3.2 Classifier Identifier

The value of the field specifies an identifier for the Classifier. This value is unique to per Service Flow. The CMTS assigns the Packet Classifier Identifier.

Type	Length	Value
[22/23].2	2	1 - 65535

C.2.1.3.3 Service Flow Reference

The value of the field specifies a Service Flow Reference that identifies the corresponding Service Flow.

In all Packet Classifier TLVs that occur in any message where the Service Flow ID is not known (*e.g.*, CM-initiated DSA-REQ and REG-REQ) this TLV MUST be included. In all Packet Classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ messages the Service Flow Reference MUST NOT be specified.

Type	Length	Value
[22/23].3	2	1 - 65535

C.2.1.3.4 Service Flow Identifier

The value of this field specifies the Service Flow ID that identifies the corresponding Service Flow.

In Packet Classifier TLVs where the Service Flow ID is not known, and this TLV MUST NOT be included (*e.g.*, CM-initiated DSA-REQ and REG-REQ). In Packet Classifier TLVs that occur in a DSC-REQ and CMTS-initiated DSA-REQ message, the Service Flow ID MUST be specified.

Type	Length	Value
[22/23].4	4	1 - 4,294,967,295

C.2.1.3.5 Rule Priority

The value of the field specifies the priority for the Classifier, which is used for determining the order of the Classifier. A higher value indicates higher priority.

Classifiers that appear in Configuration files and Registration messages MAY have priorities in the range 0 - 255 with the default value 0. Classifiers that appear in DSA/DSC message MUST have priorities in the range 64-191, with the default value 64.

Type	Length	Value
[22/23].5	1	

C.2.1.3.6 Classifier Activation State

The value of this field specifies whether this classifier should become active in selecting packets for the Service Flow. An inactive Classifier is typically used with an AdmittedQoSParameterSet to ensure resources are available for later activation. The actual activation of the classifier depends both on this attribute and on the state

of its service flow. If the service flow is not active then the classifier is not used, regardless of the setting of this attribute.

Type	Length	Value
[22/23].6	1	0 — Inactive 1 — Active

The default value is 1 — activate the classifier.

C.2.1.3.7 Dynamic Service Change Action

When received in a Dynamic Service Change Request, this indicates the action that should be taken with this classifier.

Type	Length	Value
[22/23].7	1	0 — DSC Add Classifier 1 — DSC Replace Classifier 2 — DSC Delete Classifier

C.2.1.4 Classifier Error Encodings

This field defines the parameters associated with Classifier Errors.

Type	Length	Value
[22/23].8	n	

A Classifier Error Encoding consists of a single Classifier Error Parameter Set which is defined by the following individual parameters: Errored Parameter, Confirmation Code and Error Message.

The Classifier Error Encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a Classifier establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the REG-RSP, DSA-RSP or DSC-RSP MUST include one Classifier Error Encoding for at least one failed Classifier requested in the REG-REQ, DSA-REQ or DSC-REQ message. A Classifier Error Encoding for the failed Classifier MUST include the Confirmation Code and Errored Parameter and MAY include an Error Message. If some Classifier Sets are rejected but other Classifier Sets are accepted, then Classifier Error Encodings MUST be included for only the rejected Classifiers.1 On success of the entire transaction, the RSP or ACK message MUST NOT include a Classifier Error Encoding.

Multiple Classifier Error Encodings may appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple Classifier parameters may be in error. A message with even a single Classifier Error Encoding MUST NOT contain any other protocol Classifier Encodings (e.g. IP, 802.1P/Q).

A Classifier Error Encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.1.4.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Classifier parameter in error in a rejected Classifier request. A Classifier Error Parameter Set MUST have exactly one Errored Parameter TLV within a given Classifier Error Encoding.

Subtype	Length	Value
[22/23].8.1	n	Classifier Encoding Subtype in Error

If the length is one, then the value is the single-level subtype where the error was found, *e.g.*, 7 indicates an invalid Change Action. If the length is two, then the value is the multi-level subtype where there error was found *e.g.*, 9-2 indicates an invalid IP Protocol value.

C.2.1.4.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in Appendix C.4. A Classifier Error Parameter Set MUST have exactly one Error Code within a given Classifier Error Encoding.

Subtype	Length	Value
[22/23].8.2	1	Confirmation code

A value of okay(0) indicates that the Classifier request was successful. Since a Classifier Error Parameter Set is only applies to errored parameters, this value MUST NOT be used.

C.2.1.4.3 Error Message

This subtype is optional in a Classifier Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Classifier request. A Classifier Error Parameter Set MAY have zero or one Error Message subtypes within a given Classifier Error Encoding.

SubType	Length	Value
[22/23].8.3	n	Zero-terminated string of ASCII characters.

Note: The length N includes the terminating zero.

Note: The entire Classifier Encoding message MUST have a total length of less than 256 characters.

C.2.1.5 IP Packet Classification Encodings

This field defines the parameters associated with IP packet classification.

Type	Length	Value
[22/23].9	n	

C.2.1.5.1 IP Type of Service Range and Mask

The values of the field specify the matching parameters for the IP ToS byte range and mask. An IP packet with IP ToS byte value "ip-tos" matches this parameter if tos-low <= (ip-tos AND tos-mask) <= tos-high. If this field is omitted, then comparison of the IP packet ToS byte for this entry is irrelevant.

Type	Length	Value
[22/23].9.1	3	tos-low, tos-high, tos-mask

C.2.1.5.2 IP Protocol

The value of the field specifies the matching value for the IP Protocol field [RFC-1700]. If this parameter is omitted, then comparison of the IP header Protocol field for this entry is irrelevant.

There are two special IP Protocol field values: "256" matches traffic with any IP Protocol value, and "257" matches both TCP and UDP traffic. An entry that includes an IP Protocol field value greater than 257 MUST be invalidated for comparisons (*i.e.*, no traffic can match this entry).

Type	Length	Value
[22/23].9.2	2	prot1, prot2

Valid Range
0 — 257

C.2.1.5.3 IP Source Address

The value of the field specifies the matching value for the IP source address. An IP packet with IP source address "ip-src" matches this parameter if src = (ip-src AND smask), where "smask" is the parameter from Appendix C.2.1.5.4. If this parameter is omitted, then comparison of the IP packet source address for this entry is irrelevant.

Type	Length	Value
[22/23].9.3	4	src1, src2, src3, src4

C.2.1.5.4 IP Source Mask

The value of the field specifies the mask value for the IP source address, as described in Appendix C.2.1.5.3. If this parameter is omitted, then the default IP source mask is 255.255.255.255.

Type	Length	Value
[22/23].9.4	4	smask1, smask2, smask3, smask4

C.2.1.5.5 IP Destination Address

The value of the field specifies the matching value for the IP destination address. An IP packet with IP destination address "ip-dst" matches this parameter if dst = (ip-dst AND dmask), where "dmask" is the parameter from Appendix C.2.1.5.6. If this parameter is omitted, then comparison of the IP packet destination address for this entry is irrelevant.

Type	Length	Value
[22/23].9.5	4	dst1, dst2, dst3, dst4

C.2.1.5.6 IP Destination Mask

The value of the field specifies the mask value for the IP destination address, as described in Appendix C.2.1.5.5, IP Destination Address. If this parameter is omitted, then the default IP destination mask is 255.255.255.255.

Type	Length	Value
[22/23].9.6	4	dmask1, dmask2, dmask3, dmask4

C.2.1.5.7 TCP/UDP Source Port Start

The value of the field specifies the low-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow <= src-port <= sporthigh. If this parameter is omitted, then the default value of sportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.7	2	sportlow1, sportlow2

C.2.1.5.8 TCP/UDP Source Port End

The value of the field specifies the high-end TCP/UDP source port value. An IP packet with TCP/UDP port value "src-port" matches this parameter if sportlow <= src-port <= sporthigh. If this parameter is omitted, then the default value of sporthigh is 65535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.8	2	sporthigh1, sporthigh2

C.2.1.5.9 TCP/UDP Destination Port Start

The value of the field specifies the low-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow <= dst-port <=dporthigh. If this parameter is omitted, then the default value of dportlow is 0. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.9	2	dportlow1, dportlow2

C.2.1.5.10 TCP/UDP Destination Port End

The value of the field specifies the high-end TCP/UDP destination port value. An IP packet with TCP/UDP port value "dst-port" matches this parameter if dportlow <= dst-port <= dporthigh. If this parameter is omitted, then the default value of dporthigh is 65535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23].9.10	2	dporthigh1, dporthigh2

C.2.1.6 Ethernet LLC Packet Classification Encodings

This field defines the parameters associated with Ethernet LLC packet classification.

Type	Length	Value
[22/23].10	n	

C.2.1.6.1 Destination MAC Address

The values of the field specifies the matching parameters for the MAC destination address. An Ethernet packet with MAC destination address "etherdst" matches this parameter if dst = (etherdst AND msk). If this parameter is omitted, then comparison of the Ethernet MAC destination address for this entry is irrelevant.

Type	Length	Value
[22/23].10.1	12	dst1, dst2, dst3, dst4, dst5, dst6, msk1, msk2, msk3, msk4, msk5, msk6

C.2.1.6.2 Source MAC Address

The value of the field specifies the matching value for the MAC source address. If this parameter is omitted, then comparison of the Ethernet MAC source address for this entry is irrelevant.

Type	Length	Value
[22/23].10.2	6	src1, src2, src3, src4, src5, src6

C.2.1.6.3 Ethertype/DSAP/MacType

type, eprot1, and eprot2 indicate the format of the layer 3 protocol ID in the Ethernet packet as follows:

If type = 0, the rule does not use the layer 3 protocol type as a matching criteria. If type = 0, eprot1, eprot2 are ignored when considering whether a packet matches the current rule.

If type = 1, the rule applies only to frames which contain an Ethertype value. Ethertype values are contained in packets using the DEC-Intel-Xerox (DIX) encapsulation or the [RFC-1042] Sub-Network Access Protocol (SNAP) encapsulation formats. If type = 1, then eprot1, eprot2 gives the 16-bit value of the Ethertype that the packet must match in order to match the rule

If type = 2, the rule applies only to frames using the IEEE 802.2 encapsulation format with a Destination Service (DSAP) other than 0xAA (which is reserved for SNAP). If type = 2, the lower 8 bits of the eprot1, eprot2, MUST match the DSAP byte of the packet in order to match the rule.

If type = 3, the rule applies only to MAC Management Messages (FC field 1100001x) with a "type" field of its MAC Management Message header (6.3.1) between the values of eprot1 and eprot2, inclusive. As exceptions, the following MAC Management message types MUST NOT be classified, and are always transmitted on the primary service flow:

Type 4: RNG_REQ
Type 6: REG_REQ
Type 7: REG_RSP
Type 14: REG_ACK

If type = 4, the rule is considered a "catch-all" rule that matches all Data PDU packets. The rule does not match MAC Management Messages. The value of eprot1 and eprot2 are ignored in this case.

If the Ethernet frame contains an 802.1P/Q Tag header (*i.e.*, Ethertype 0x8100), this object applies to the embedded Ethertype field within the 802.1P/Q header.

Other values of type are reserved. If this TLV is omitted, then comparison of either the Ethertype or IEEE 802.2 DSAP for this rule is irrelevant.

Type	Length	Value
[22/23].10.3	3	type, eprot1, eprot2

C.2.1.7 IEEE 802.1P/Q Packet Classification Encodings

This field defines the parameters associated with IEEE 802.1P/Q packet classification.

Type	Length	Value
[22/23].11	n	

C.2.1.7.1 IEEE 802.1P User_Priority

The values of the field specify the matching parameters for the IEEE 802.1P user_priority bits. An Ethernet packet with IEEE 802.1P user_priority value "priority" matches these parameters if pri-low <= priority <= pri-high. If this field is omitted, then comparison of the IEEE 802.1P user_priority bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation MUST NOT match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q encapsulated traffic, then this entry MUST NOT be used for any traffic.

Type	Length	Value
[22/23].11.1	2	pri-low, pri-high

Valid Range
0 — 7 for pri-low and pri-high

C.2.1.7.2 IEEE 802.1Q VLAN_ID

The value of the field specify the matching value for the IEEE 802.1Q vlan_id bits. Only the first (*i.e.*, most-significant) 12 bits of the specified vlan_id field are significant; the final four bits MUST be ignored for comparison. If this field is omitted, then comparison of the IEEE 802.1Q vlan_id bits for this entry is irrelevant.

If this parameter is specified for an entry, then Ethernet packets without IEEE 802.1Q encapsulation MUST NOT match this entry. If this parameter is specified for an entry on a CM that does not support forwarding of IEEE 802.1Q encapsulated traffic, then this entry MUST NOT be used for any traffic.

Type	Length	Value
[22/23].11.2	2	vlan_id1, vlan_id2

C.2.1.7.3 Vendor Specific Classifier Parameters

This allows vendors to encode vendor-specific Classifier parameters. The Vendor ID MUST be the first TLV embedded inside Vendor Specific Classifier Parameters. If the first TLV inside Vendor Specific Classifier Parameters is not a Vendor ID, then the TLV MUST be discarded. (Refer to Appendix C.1.1.17)

Type	Length	Value
[22/23].43	n	

C.2.2 Service Flow Encodings

The following type/length/value encodings MUST be used in the configuration file, registration messages, and Dynamic Service messages to encode parameters for Service Flows. 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.2.2.1 Upstream Service Flow Encodings

This field defines the parameters associated with upstream scheduling for a Service Flow. It is somewhat complex in that is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream Service Flow configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings. These type fields are not valid in other encoding contexts.

Type	Length	Value
24	n	

C.2.2.2 Downstream Service Flow Encodings

This field defines the parameters associated with downstream scheduling for a Service Flow. It is somewhat complex in that is composed from a number of encapsulated type/length/value fields.

Note that the encapsulated upstream and downstream flow classification configuration setting strings share the same subtype field numbering plan, because many of the subtype fields defined are valid for both types of configuration settings except Service Flow encodings. These type fields are not valid in other encoding contexts.

Type	Length	Value
25	n	

C.2.2.3 General Service Flow Encodings

C.2.2.3.1 Service Flow Reference

The Service Flow Reference is used to associate a packet classifier encoding with a Service Flow encoding. A Service Flow Reference is only used to establish a Service Flow ID. Once the Service Flow exists and has an assigned Service Flow ID, the Service Flow Reference MUST no longer be used. The Service Flow Reference is unique per configuration file, Registration message exchange, or Dynamic Service Add message exchange.

Type	Length	Value
[24/25].1	2	1 - 65535

C.2.2.3.2 Service Flow Identifier

The Service Flow Identifier is used by the CMTS as the primary reference of a Service Flow. Only the CMTS can issue a Service Flow Identifier. It uses this parameterization to issue Service Flow Identifiers in CMTS-initiated DSA-Requests and in its REG/DSA-Response to CM-initiated REG/DSA-Requests. The CM specifies the SFID of a service flow using this parameter in a DSC-REQ message. Both the CM and CMTS MAY use this TLV to encode Service Flow IDs in a DSD-REQ.

The configuration file MUST NOT contain this parameter.

Type	Length	Value
[24/25].2	4	1 - 4,294,967,295

C.2.2.3.3 Service Identifier

The value of this field specifies the Service Identifier assigned by the CMTS to a Service Flow with a non-null AdmittedQosParameterSet or ActiveQosParameterSet. This is used in the bandwidth allocation MAP to assign upstream bandwidth. This field MUST be present in CMTS-initiated DSA-REQ or DSC-REQ message related to establishing an admitted or active upstream Service Flow. This field MUST also be present in REG-RSP, DSA-RSP, and DSC-RSP messages related to the successful establishment of an admitted or active upstream Service Flow. This field MUST NOT be present in settings related to downstream Service Flows; the Service Identifier only applies to upstream Service Flows.

Even though a Service Flow has been successfully admitted or activated (i.e., has an assigned Service ID) the Service Flow ID MUST be used for subsequent DSx message signaling as it is the primary handle for a service flow. If a Service Flow is no longer admitted or active (via DSC-REQ) its Service ID MAY be reassigned by the CMTS.

SubType	Length	Value
[24].3	2	SID (low-order 14 bits)

C.2.2.3.4 Service Class Name

The value of the field refers to a predefined CMTS service configuration to be used for this Service Flow.

Type	Length	Value
[24/25].4	2 to 16	Zero-terminated string of ASCII characters.

Note: The length includes the terminating zero.

When the Service Class Name is used in a Service Flow encoding, it indicates that all the unspecified QoS Parameters of the Service Flow need to be provided by the CMTS. It is up to the operator to synchronize the definition of Service Class Names in the CMTS and in the configuration file.

C.2.2.3.5 Quality of Service Parameter Set Type

This parameter MUST appear within every Service Flow Encoding, with the exception of Service Flow Encodings in the DSD-REQ where the Quality of Service Parameter Set Type has no value. It specifies the

proper application of the QoS Parameter Set or Service Class Name: to the Provisioned set, the Admitted set, and/or the Active set. When two QoS Parameter Sets are the same, a multi-bit value of this parameter MAY be used to apply the QoS parameters to more than one set. A single message MAY contain multiple QoS parameter sets in separate type 24/25 Service Flow Encodings for the same Service Flow. This allows specification of the QoS Parameter Sets when their parameters are different. Bit 0 is the LSB of the Value field.

For every Service Flow that appears in a Registration-Request or Registration-Response message, there MUST be a Service Flow Encoding that specifies a ProvisionedQoSParameterSet. This Service Flow Encoding, or other Service Flow Encoding(s), MAY also specify an Admitted and/or Active set.

Any Service Flow Encoding that appears in a Dynamic Service Message MUST NOT specify the ProvisionedQoSParameterSet.

Type	Length	Value	
[24/25].6	1	Bit # 0	Provisioned Set
		Bit # 1	Admitted Set
		Bit # 2	Active Set

Table C-2. Values Used in REG-REQ and REG-RSP Messages

Value	Messages
001	Apply to Provisioned set only
011	Apply to Provisioned and Admitted set, and perform admission control
101	Apply to Provisioned and Active sets, perform admission control on Admitted set in separate Service Flow Encoding, and activate the Service flow.
111	Apply to Provisioned, Admitted, and Active sets; perform admission control and activate this Service Flow

Table C-3. Values Used In REG-REQ, REG-RSP, and Dynamic Service Messages

Value	Messages
010	Perform admission control and apply to Admitted set
100	Check against Admitted set in separate Service flow Encoding, perform admission control if needed, activate this Service Flow, and apply to Active set
110	Perform admission control and activate this Service Flow, apply parameters to both Admitted and Active sets

The value 000 is used only in Dynamic Service Change messages. It is used to set the Active and Admitted sets to Null (see Section 8.1.7.4).

A CMTS MUST handle a single update to each of the Active and Admitted QoS parameter sets. The ability to process multiple Service Flow Encodings that specify the same QoS parameter set is NOT required, and is left as a vendor-specific function. If a DSA/DSC contains multiple updates to a single QoS parameter set and the vendor does not support such updates, then the CMTS MUST reply with error code 2, reject-unrecognized-configuration-setting.

C.2.2.4 Service Flow Error Encodings

This field defines the parameters associated with Service Flow Errors.

Type	Length	Value
[24/25].5	n	

A Service Flow Error Encoding consists of a single Service Flow Error Parameter Set which is defined by the following individual parameters: Errored Parameter, Confirmation Code and Error Message.

The Service Flow Error Encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a Service Flow establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

The Service Flow Error Encoding is returned in REG-ACK, DSA-ACK and DSC-ACK messages to indicate the reason for the recipient's negative response to the expansion of a Service Class Name in a corresponding REG-RSP, DSA-RSP or DSC-RSP.

On failure, the REG-RSP, DSA-RSP or DSC-RSP MUST include one Service Flow Error Encoding for at least one failed Service Flow requested in the REG-REQ, DSA-REQ or DSC-REQ message. On failure, the REG-ACK, DSA-ACK or DSC-ACK MUST include one Service Flow Error Encoding for at least one failed Service Class Name expansion in the REG-RSP, DSA-RSP or DSC-RSP message. A Service Flow Error Encoding for the failed Service Flow MUST include the Confirmation Code and Errored Parameter and MAY include an Error Message. If some Service Flow Parameter Sets are rejected but other Service Flow Parameter Sets are accepted, then Service Flow Error Encodings MUST be included for only the rejected Service Flow.

On success of the entire transaction, the RSP or ACK message MUST NOT include a Service Flow Error Encoding.

Multiple Service Flow Error Encodings MAY appear in a REG-RSP, DSA-RSP, DSC-RSP, REG-ACK, DSA-ACK or DSC-ACK message, since multiple Service Flow parameters may be in error. A message with even a single Service Flow Error Encoding MUST NOT contain any QoS Parameters.

A Service Flow Error Encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.4.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Service Flow parameter in error in a rejected Service Flow request or Service Class Name expansion response.A Service Flow Error Parameter Set MUST have exactly one Errored Parameter TLV within a given Service Flow Error Encoding.

Subtype	Length	Value
[24/25].5.1	1	Service Flow Encoding Subtype in Error

C.2.2.4.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in Appendix C.4. A Service Flow Error Parameter Set MUST have exactly one Error Code within a given Service Flow Error Encoding.

Subtype	Length	Value
[24/25].5.2	1	Confirmation code

A value of okay(0) indicates that the Service Flow request was successful. Since a Service Flow Error Parameter Set only applies to errored parameters, this value MUST NOT be used.

C.2.2.4.3 Error Message

This subtype is optional in a Service Flow Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Service Flow request. A Service Flow Error Parameter Set MAY have zero or one Error Message subtypes within a given Service Flow Error Encoding.

SubType	Length	Value
[24/25].5.3	n	Zero-terminated string of ASCII characters.

Note: The length N includes the terminating zero.

Note: The entire Service Flow Encoding message MUST have a total length of less than 256 characters.

C.2.2.5 Common Upstream and Downstream Quality-of-Service Parameter Encodings

The remaining Type 24 & 25 parameters are QoS Parameters. Any given QoS Parameter type MUST appear zero or one times per Service Flow Encoding.

C.2.2.5.1 Traffic Priority

The value of this parameter specifies the priority assigned to a Service Flow. Given two Service Flows identical in all QoS parameters besides priority, the higher priority Service Flow SHOULD be given lower delay and higher buffering preference. For otherwise non-identical Service Flows, the priority parameter SHOULD NOT take precedence over any conflicting Service Flow QoS parameter. The specific algorithm for enforcing this parameter is not mandated here.

For upstream service flows, the CMTS SHOULD use this parameter when determining precedence in request service and grant generation, and the CM MUST preferentially select contention Request opportunities for Priority Request Service IDs (refer to Appendix A.2.3) based on this priority and its Request/Transmission Policy (refer to Appendix C.2.2.6.3).

Type	Length	Value
[24/25].7	1	0 to 7 — Higher numbers indicate higher priority

Note: The default priority is 0.

C.2.2.5.2 Maximum Sustained Traffic Rate

This parameter is the rate parameter R of a token-bucket-based rate limit for packets. R is expressed in bits per second, and MUST take into account all MAC frame data PDU of the Service Flow from the byte following the MAC header HCS to the end of the CRC[1]. The number of bytes forwarded (in bytes) is limited during any time interval T by Max(T), as described in the expression

$$Max(T) = T * (R / 8) + B, \quad (1)$$

where the parameter B (in bytes) is the Maximum Traffic Burst Configuration Setting (refer to Appendix C.2.2.5.3).

1. The payload size includes every PDU in a Concatenated MAC Frame.

Note: This parameter does not limit the instantaneous rate of the Service Flow.

Note: The specific algorithm for enforcing this parameter is not mandated here. Any implementation which satisfies the above equation is conformant.

Note: If this parameter is omitted or set to zero, then there is no explicitly-enforced traffic rate maximum. This field specifies only a bound, not a guarantee that this rate is available.

C.2.2.5.2.1 Upstream Maximum Sustained Traffic Rate

For an upstream Service Flow, the CM MUST NOT request bandwidth exceeding the Max(T) requirement in (1) during any interval T because this could force the CMTS to fill MAPs with deferred grants.

The CM MUST defer upstream packets that violate (1) and "rate shape" them to meet the expression, up to a limit as implemented by vendor buffering restrictions.

The CMTS MUST enforce expression (1) on all upstream data transmissions, including data sent in contention. The CMTS MAY consider unused grants in calculations involving this parameter. The CMTS MAY 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, or (c) discarding over-limit data packets. A CMTS MUST report this condition to a policy module. If the CMTS is policing by discarding either packets or requests, the CMTS MUST allow a margin of error between the CM and CMTS algorithms.

Type	Length	Value
24.8	4	R (in bits per second)

C.2.2.5.2.2 Downstream Maximum Sustained Traffic Rate

For a downstream Service Flow, this parameter is only applicable at the CMTS. The CMTS MUST enforce expression (1) on all downstream data transmissions. The CMTS MUST NOT forward downstream packets that violates (1) in any interval T. The CMTS SHOULD "rate shape" the downstream traffic by enqueuing packets arriving in excess of (1), and delay them until the expression can be met.

This parameter is not intended for enforcement on the CM.

Type	Length	Value
25.8	4	R (in bits per second)

C.2.2.5.3 Maximum Traffic Burst

The value of this parameter specifies the token bucket size B (in bytes) for this Service Flow as described in expression (1). This value is calculated from the byte following the MAC header HCS to the end of the CRC[2].

The minimum value of B is 1522 bytes. If this parameter is omitted, the default value for B is 3044 bytes. This parameter has no effect unless a non-zero value has been provided for the Maximum Sustained Traffic Rate parameter.

2. The payload size includes every PDU in a Concatenated MAC Frame.

For an upstream service flow, if B is sufficiently less than the Maximum Concatenated Burst parameter, then enforcement of the rate limit equation will limit the maximum size of a concatenated burst.

Type	Length	Value
[24/25].9	4	B (bytes)

Note: The specific algorithm for enforcing this parameter is not mandated here. Any implementation which satisfies the above equation is conformant.

Note: The value of this parameter affects the tradeoff between the data latency perceived by an individual application, and the traffic engineering requirements of the network. A large value will tend to reduce the latency introduced by rate limiting for applications with bursty traffic patterns. A small value will tend to spread out the bursts of data generated by such applications, which may benefit traffic engineering within the network.

C.2.2.5.4 Minimum Reserved Traffic Rate

This parameter specifies the minimum rate, in bits/sec, reserved for this Service Flow. The CMTS SHOULD be able to satisfy bandwidth requests for a Service Flow up to its Minimum Reserved Traffic Rate. If less bandwidth than its Minimum Reserved Traffic Rate is requested for a Service Flow, the CMTS MAY reallocate the excess reserved bandwidth for other purposes. The aggregate Minimum Reserved Traffic Rate of all Service Flows MAY exceed the amount of available bandwidth. This value of this parameter is calculated from the byte following the MAC header HCS to the end of the CRC[3]. If this parameter is omitted, then it defaults to a value of 0 bits/sec (*i.e.*, no bandwidth is reserved for the flow by default).

This field is only applicable at the CMTS and MUST be enforced by the CMTS.

Type	Length	Value
[24/25].10	4	

Note: The specific algorithm for enforcing the value specified in this field is not mandated here.

C.2.2.5.5 Assumed Minimum Reserved Rate Packet Size

The value of this field specifies an assumed minimum packet size (in bytes) for which the Minimum Reserved Traffic Rate will be provided. This parameter is defined in bytes and is specified as the bytes following the MAC header HCS to the end of the CRC[4]. If the Service Flow sends packets of a size smaller than this specified value, such packets will be treated as being of the size specified in this parameter for calculating the minimum Reserved Traffic Rate and for calculating bytes counts (*e.g.*, bytes transmitted) which may ultimately be used for billing.

The CMTS MUST apply this parameter to its Minimum Reserved Traffic Rate algorithm. This parameter is used by the CMTS to estimate the per packet overhead of each packet in the service flow.

If this parameter is omitted, then the default value is CMTS implementation dependent.

Type	Length	Value
[24/25].11	2	

3. The payload size includes every PDU in a Concatenated MAC Frame.

4. The payload size includes every PDU in a Concatenated MAC Frame.

C.2.2.5.6 Timeout for Active QoS Parameters

The value of this parameter specifies the maximum duration resources remain unused on an active Service Flow. If there is no activity on the Service Flow within this time interval, the CMTS MUST change the active and admitted QoS Parameter Sets to null. The CMTS MUST signal this resource change with a DSC-REQ to the CM.

Type	Length	Value
[24/25].12	2	seconds

This parameter MUST be enforced at the CMTS and SHOULD NOT be enforced at the CM. The parameter is processed by the CMTS for every QoS set contained in Registration messages and Dynamic Service messages. If the parameter is omitted, the default of 0 (*i.e.*, infinite timeout) is assumed. The value specified for the active QoS set must be less than or equal to the corresponding value in the admitted QoS set which must be less than or equal to the corresponding value in the provisioned/authorized QoS set. If the requested value is too large, the CMTS MAY reject the message or respond with a value less than that requested. If the Registration or Dynamic Service message is accepted by the CMTS and acknowledged by the CM, the Active MQoS Timeout timer is loaded with the new value of the timeout. The timer is activated if the message activates the associated Service Flow. The timer is deactivated if the message sets the active QoS set to null.

C.2.2.5.7 Timeout for Admitted QoS Parameters

The value of this parameter specifies the duration that the CMTS MUST hold resources for a Service Flow's Admitted QoS Parameter Set while they are in excess of its Active QoS Parameter Set. If there is no DSC-REQ to activate the Admitted QoS Parameter Set within this time interval, and there is no DSC to refresh the QoS parameter sets and restart the timeout (see Section 8.1.5.2), the resources that are admitted but not activated MUST be released, and only the active resources retained. The CMTS MUST set the Admitted QoS Parameter Set equal to the Active QoS Parameter Set for the Service Flow and initiate a DSC-REQ exchange with the CM to inform it of the change.

Type	Length	Value
[24/25].13	2	seconds

This parameter MUST be enforced at the CMTS and SHOULD NOT be enforced at the CM. The parameter is processed by the CMTS for every QoS set contained in Registration messages and Dynamic Service messages. If the parameter is omitted, the default of 200 seconds is assumed. A value of 0 means that the Service Flow can remain in the admitted state for an infinite amount of time and MUST NOT be timed out due to inactivity. However, this is subject to policy control by the CMTS. The value specified for the active QoS set must be less than or equal to the corresponding value in the admitted QoS set which must be less than or equal to the corresponding value in the provisioned/authorized QoS set. If the requested value is too large, the CMTS MAY reject the message or respond with a value less than that requested. If the Registration or Dynamic Service message containing this parameter is accepted by the CMTS and acknowledged by the CM, the Admitted QoS Timeout timer is loaded with the new value of the timeout. The timer is activated if the message admits resources greater than the active set. The timer is deactivated if the message sets the active QoS set and admitted QoS set equal to each other.

C.2.2.5.8 Vendor Specific QoS Parameters

This allows vendors to encode vendor-specific QoS parameters. The Vendor ID MUST be the first TLV embedded inside Vendor Specific QoS Parameters. If the first TLV inside Vendor Specific QoS Parameters is not a Vendor ID, then the TLV MUST be discarded. (Refer to Appendix C.1.1.17)

Type	Length	Value
[24/25].43	n	

C.2.2.6 Upstream-Specific QoS Parameter Encodings

C.2.2.6.1 Maximum Concatenated Burst

The value of this parameter specifies the maximum concatenated burst (in bytes) which a Service Flow is allowed. This parameter is calculated from the FC byte of the Concatenation MAC Header to the last CRC in the concatenated MAC frame.

A value of 0 means there is no limit. If this parameter is omitted the default value is 1522.

This field is only applicable at the CM. If defined, this parameter MUST be enforced at the CM.

Note: This value does not include any physical layer overhead.

Type	Length	Value
24.14	2	

Note: This applies only to concatenated bursts. It is legal and, in fact, it may be useful to set this smaller than the maximum Ethernet packet size. Of course, it is also legal to set this equal to or larger than the maximum Ethernet packet size.

Note: The maximum size of a concatenated burst can also be limited by the enforcement of a rate limit, if the Maximum Traffic Burst parameter is small enough, and by limits on the size of data grants in the UCD message.

C.2.2.6.2 Service Flow Scheduling Type

The value of this parameter specifies which upstream scheduling service is used for upstream transmission requests and packet transmissions. If this parameter is omitted, then the Best Effort service MUST be assumed.

This parameter is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.15	1	0 Reserved 1 for Undefined (CMTS implementation-dependent[5]) 2 for Best Effort 3 for Non-Real-Time Polling Service 4 for Real-Time Polling Service 5 for Unsolicited Grant Service with Activity Detection 6 for Unsolicited Grant Service 7 through 255 are reserved for future use

5. The specific implementation dependent scheduling service type could be defined in the 24.43 Vendor Specific Information Field.

C.2.2.6.3 Request/Transmission Policy

The value of this parameter specifies which IUC opportunities the CM uses for upstream transmission requests and packet transmissions for this Service Flow, whether requests for this Service Flow may be piggybacked with data and whether data packets transmitted on this Service Flow can be concatenated, fragmented, or have their payload headers suppressed. For UGS, it also specifies how to treat packets that do not fit into the UGS grant. See Section 8.2 for requirements related to settings of the bits of this parameter for each Service Flow Scheduling Type.

This parameter is required for all Service Flow Scheduling Types except Best Effort. If omitted in a Best Effort Service Flow QoS parameter Set, the default value of zero MUST be used. Bit #0 is the LSB of the Value field. Bits are set to 1 to select the behavior defined below:

Type	Length	Value
24.16	4	Bit #0 The Service Flow MUST NOT use “all CMs” broadcast request opportunities. Bit #1 The Service Flow MUST NOT use Priority Request multicast request opportunities. (Refer to Appendix A.2.3) Bit #2 The Service Flow MUST NOT use Request/Data opportunities for Requests Bit #3 The Service Flow MUST NOT use Request/Data opportunities for Data Bit #4 The Service Flow MUST NOT piggyback requests with data. Bit #5 The Service Flow MUST NOT concatenate data. Bit #6 The Service Flow MUST NOT fragment data Bit #7 The Service Flow MUST NOT suppress payload headers Bit #8[6] The Service Flow MUST drop packets that do not fit in the Unsolicited Grant Size[7] All other bits are reserved.

Note: Data grants include both short and long data grants.

C.2.2.6.4 Nominal Polling Interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive unicast request opportunities for this Service Flow on the upstream channel. This parameter is typically suited for Real-Time and Non-Real-Time Polling Service.

The ideal schedule for enforcing this parameter is defined by a reference time t_0, with the desired transmission times $t_i = t_0 + i*\text{interval}$. The actual poll times, t'_i MUST be in the range $t_i <= t'_i <= t_i + \text{jitter}$, where interval is the value specified with this TLV, and jitter is Tolerated Poll Jitter. The accuracy of the ideal poll times, t_i, are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.17	4	µsec

6. This bit only applies to Service Flows with the Unsolicited Grant Service Flow Scheduling Type, if this bit is set on any other Service Flow Scheduling type it MUST be ignored

7. Packets that classify to an Unsolicited Grant Service Flow and are larger than the Grant Size associated with that Service Flow are normally transmitted on the Primary Service Flow. This parameter overrides that default behavior.

C.2.2.6.5 Tolerated Poll Jitter

The values in this parameter specifies the maximum amount of time that the unicast request interval may be delayed from the nominal periodic schedule (measured in microseconds) for this Service Flow.

The ideal schedule for enforcing this parameter is defined by a reference time t_0, with the desired poll times $t_i = t_0 + i*interval$. The actual poll, t'_i MUST be in the range $t_i <= t'_i <= t_i + jitter$, where jitter is the value specified with this TLV and interval is the Nominal Poll Interval. The accuracy of the ideal poll times, t_i, are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This parameter is only applicable at the CMTS. If defined, this parameter represents a service commitment (or admission criteria) at the CMTS.

Type	Length	Value
24.18	4	µsec

C.2.2.6.6 Unsolicited Grant Size

The value of this parameter specifies the unsolicited grant size in bytes. The grant size includes the entire MAC frame data PDU from the Frame Control byte to end of the MAC frame.

This parameter is applicable at the CMTS and MUST be enforced at the CMTS.

Type	Length	Value
24.19	2	

Note: For UGS, this parameter should be used by the CMTS to compute the size of the unsolicited grant in minislots.

C.2.2.6.7 Nominal Grant Interval

The value of this parameter specifies the nominal interval (in units of microseconds) between successive data grant opportunities for this Service Flow. This parameter is required for Unsolicited Grant and Unsolicited Grant with Activity Detection Service Flows.

The ideal schedule for enforcing this parameter is defined by a reference time t_0, with the desired transmission times $t_i = t_0 + i*interval$. The actual grant times, t'_i MUST be in the range $t_i <= t'_i <= t_i + jitter$, where interval is the value specified with this TLV, and jitter is the Tolerated Grant Jitter. When multiple grants per interval are requested, all grants MUST be within this interval, thus the Nominal Grant Interval and Tolerated Grant Jitter MUST be maintained by the CMTS for all grants in this Service Flow. The accuracy of the ideal grant times, t_i, are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.20	4	µsec

C.2.2.6.8 Tolerated Grant Jitter

The values in this parameter specifies the maximum amount of time that the transmission opportunities may be delayed from the nominal periodic schedule (measured in microseconds) for this Service Flow.

The ideal schedule for enforcing this parameter is defined by a reference time t_0, with the desired transmission times $t_i = t_0 + i*interval$. The actual transmission opportunities, t'_i MUST be in the range $t_i <= t'_i <= t_i + jitter$, where jitter is the value specified with this TLV and interval is the Nominal Grant Interval. The accuracy of the ideal grant times, t_i, are measured relative to the CMTS Master Clock used to generate timestamps (refer to Section 7.3).

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value
24.21	4	µsec

C.2.2.6.9 Grants per Interval

For Unsolicited Grant Service, the value of this parameter indicates the actual number of data grants per Nominal Grant Interval. For Unsolicited Grant Service with Activity Detection, the value of this parameter indicates the maximum number of Active Grants per Nominal Grant Interval. This is intended to enable the addition of sessions to an existing Unsolicited Grant Service Flow via the Dynamic Service Change mechanism, without negatively impacting existing sessions.

The ideal schedule for enforcing this parameter is defined by a reference time t_0, with the desired transmission times $t_i = t_0 + i*interval$. The actual grant times, t'_i MUST be in the range $t_i <= t'_i <= t_i + jitter$, where interval is the Nominal Grant Interval, and jitter is the Tolerated Grant Jitter. When multiple grants per interval are requested, all grants MUST be within this interval, thus the Nominal Grant Interval and Tolerated Grant Jitter MUST be maintained by the CMTS for all grants in this Service Flow.

This field is mandatory for Unsolicited Grant and Unsolicited Grant with Activity Detection Scheduling Types. This field is only applicable at the CMTS, and MUST be enforced by the CMTS.

Type	Length	Value	Valid Range
24.22	1	# of grants	0-127

C.2.2.6.10 IP Type of Service Overwrite

The CMTS MUST overwrite IP packets with IP ToS byte value “orig-ip-tos” with the value “new-ip-tos”, where new-ip-tos = ((orig-ip-tos AND tos-and-mask) OR tos-or-mask). If this parameter is omitted, then the IP packet ToS byte is not overwritten.

This parameter is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.23	2	tos-and-mask, tos-or-mask

C.2.2.6.11 Unsolicited Grant Time Reference

For Unsolicited Grant Service and Unsolicited Grant Service with Activity Detection, the value of this parameter specifies a reference time t_0 from which can be derived the desired transmission times $t_i = t_0 + i*interval$, where interval is the Nominal Grant Interval (Refer to Appendix C.2.2.6.7). This parameter is applicable only for messages transmitted from the CMTS to the CM, and only when a UGS or UGS-AD service flow is being made active. In such cases this is a mandatory parameter.

Type	Length	Value	Valid Range
24.24	4	CMTS Timestamp	0-4,294,967,295

The timestamp specified in this parameter represents a count state of the CMTS 10.24 MHz master clock. Since a UGS or UGS-AD service flow is always activated before transmission of this parameter to the modem, the reference time t_0 is to be interpreted by the modem as the ideal time of the next grant only if t_0 follows the current time. If t_0 precedes the current time, the modem can calculate the offset from the current time to the ideal time of the next grant according to:

interval - (((current time - t_0)/10.24) modulus interval)
where: interval is in units of microseconds
current time and t_0 are in 10.24 MHz units

C.2.2.7 Downstream-Specific QoS Parameter Encodings

C.2.2.7.1 Maximum Downstream Latency

The value of this parameter specifies the maximum latency between the reception of a packet by the CMTS on its NSI and the forwarding of the packet to its RF Interface.

If defined, this parameter represents a service commitment (or admission criteria) at the CMTS and MUST be guaranteed by the CMTS. A CMTS does not have to meet this service commitment for Service Flows that exceed their minimum downstream reserved rate.

Type	Length	Value
25.14	4	µsec

C.2.2.8 Payload Header Suppression

This field defines the parameters associated with Payload Header Suppression.

Type	Length	Value
26	n	

Note: The entire Payload Header Suppression TLV MUST have a length of less than 255 characters.

C.2.2.8.1 Classifier Reference

The value of the field specifies a Classifier Reference that identifies the corresponding Classifier. (Refer to Appendix C.2.1.3.1)

Type	Length	Value
26.1	1	1 - 255

C.2.2.8.2 Classifier Identifier

The value of the field specifies a Classifier Identifier that identifies the corresponding Classifier. (Refer to Appendix C.2.1.3.2)

Type	Length	Value
26.2	2	1 - 65535

C.2.2.8.3 Service Flow Reference

The value of the field specifies a Service Flow Reference that identifies the corresponding Service Flow. (Refer to Appendix C.2.2.3.1)

Type	Length	Value
26.3	2	1 - 65535

C.2.2.8.4 Service Flow Identifier

The value of this field specifies the Service Flow Identifier that identifies the Service Flow to which the PHS rule applies.

Type	Length	Value
26.4	4	1 - 4,294,967,295

C.2.2.8.5 Dynamic Service Change Action

When received in a Dynamic Service Change Request, this indicates the action that MUST be taken with this payload header suppression byte string.

Type	Length	Value
26.5	1	0 — Add PHS Rule 1 — Set PHS Rule 2 — Delete PHS Rule 3 — Delete all PHS Rules

The “Set PHS Rule” command is used to add specific TLVs to a partially defined payload header suppression rule. A PHS rule is partially defined when the PHSF and PHSS values are not both known. A PHS rule becomes fully defined when the PHSF and PHSS values are both known. Once a PHS rule is fully defined, “Set PHS Rule” MUST NOT be used to modify existing TLVs.

The “Delete al l PHS Rules” command is used to delete al l PHS Rules for a specified Service Flow. See Section 6.3.15 for details on DSC-REQ require d PHS parameters when using this option.

Note: An attempt to Add a PHS Rule which already exists is an error condition.

C.2.2.9 Payload Header Suppression Error Encodings

This field defines the parameters associated with Payload Header Suppression Errors.

Type	Length	Value
26.6	n	

A Payload Header Suppression Error Encoding consists of a single Payload Header Suppression Error Parameter Set which is defined by the following individual parameters: Errored Parameter, Confirmation Code and Error Message.

The Payload Header Suppression Error Encoding is returned in REG-RSP, DSA-RSP and DSC-RSP messages to indicate the reason for the recipient's negative response to a Payload Header Suppression Rule establishment request in a REG-REQ, DSA-REQ or DSC-REQ message.

On failure, the REG-RSP, DSA-RSP, or DSC-RSP MUST include one Payload Header Suppression Error Encoding for at least one failed Payload Header Suppression Rule requested in the REG-REQ, DSA-REQ or DSC-REQ message. A Payload Header Suppression Error Encoding for the failed Payload Header Suppression Rule MUST include the Confirmation Code and Errored Parameter and MAY include an Error Message. If some Payload Header Suppression Rule Sets are rejected but other Payload Header Suppression Rule Sets are accepted, then Payload Header Suppression Error Encodings MUST be included for only the rejected Payload Header Suppression Rules. On success of the entire transaction, the RSP or ACK message MUST NOT include a Payload Header Suppression Error Encoding.

Multiple Payload Header Suppression Error Encodings MAY appear in a REG-RSP, DSA-RSP or DSC-RSP message, since multiple Payload Header Suppression parameters may be in error. A message with even a single Payload Header Suppression Error Encoding MUST NOT contain any other protocol Payload Header Suppression Encodings (*e.g.*, IP, 802.1P/Q).

A Payload Header Suppression Error Encoding MUST NOT appear in any REG-REQ, DSA-REQ or DSC-REQ messages.

C.2.2.9.1 Errored Parameter

The value of this parameter identifies the subtype of a requested Payload Header Suppression parameter in error in a rejected Payload Header Suppression request. A Payload Header Suppression Error Parameter Set MUST have exactly one Errored Parameter TLV within a given Payload Header Suppression Error Encoding.

Subtype	Length	Value
26.6.1	1	Payload Header Suppression Encoding Subtype in Error

C.2.2.9.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in Appendix C.4. A Payload Header Suppression Error Parameter Set MUST have exactly one Error Code within a given Payload Header Suppression Error Encoding.

Subtype	Length	Value
26.6.2	1	Confirmation code

A value of okay(0) indicates that the Payload Header Suppression request was successful. Since a Payload Header Suppression Error Parameter Set only applies to errored parameters, this value MUST NOT be used.

C.2.2.9.3 Error Message

This subtype is optional in a Payload Header Suppression Error Parameter Set. If present, it indicates a text string to be displayed on the CM console and/or log that further describes a rejected Payload Header Suppression

request. A Payload Header Suppression Error Parameter Set MAY have zero or one Error Message subtypes within a given Payload Header Suppression Error Encoding.

SubType	Length	Value
26.6.3	n	Zero-terminated string of ASCII characters.

Note: The length n includes the terminating zero.

Note: The entire Payload Header Suppression Encoding message MUST have a total length of less than 256 characters.

C.2.2.10 Payload Header Suppression Rule Encodings

C.2.2.10.1 Payload Header Suppression Field (PHSF)

The value of this field are the bytes of the headers which MUST be suppressed by the sending entity, and MUST be restored by the receiving entity. In the upstream, the PHSF corresponds to the string of PDU bytes starting with the first byte after the MAC Header Checksum. For the downstream, the PHSF corresponds to the string of PDU bytes starting with the 13th byte after the MAC Header Checksum. This string of bytes is inclusive of both suppressed and unsuppressed bytes of the PDU header. The value of the unsuppressed bytes within the PHSF is implentation dependent.

The ordering of the bytes in the value field of the PHSF TLV string MUST follow the sequence:

Upstream
MSB of PHSF value = 1st byte of PDU
2nd MSB of PHSF value = 2nd byte of PDU
...
nth byte of PHSF (LSB of PHSF value) = nth byte of PDU

Downstream
MSB of PHSF value = 13th byte of PDU
2nd MSB of PHSF value = 14th byte of PDU
...
nth byte of PHSF (LSB of PHSF value) = (n+13)th byte of PDU

Type	Length	Value
26.7	n	string of bytes suppressed

The length n MUST always be the same as the value for PHSS.

C.2.2.10.2 Payload Header Suppression Index (PHSI)

The Payload Header Suppression Index (PHSI) has a value between 1 and 255 which uniquely references the suppressed byte string. The Index is unique per Service Flow in the upstream direction and unique per CM in the downstream direction. The upstream and downstream PHSI values are independent of each other.

Type	Length	Value
26.8	1	index value

C.2.2.10.3 Payload Header Suppression Mask (PHSM)

The value of this field is used to interpret the values in the Payload Header Suppression Field. It is used at both the sending and receiving entities on the link. The PHSM allows fields such as sequence numbers or checksums which vary in value to be excluded from suppression with the constant bytes around them suppressed.

Type	Length	Value
26.9	n	bit 0: 0 = don't suppress first byte of the suppression field 1 = suppress first byte of the suppression field bit 1: 0 = don't suppress second byte of the suppression field 1 = suppress second byte of the suppression field bit x: 0 = don't suppress (x+1) byte of the suppression field 1 = suppress (x+1) byte of the suppression field

The length n is ceiling(PHSS/8). Bit 0 is the MSB of the Value field. The value of each sequential bit in the PHSM is an attribute for the corresponding sequential byte in the PHSF.

If the bit value is a "1" (and verification passes or is disabled), the sending entity MUST suppress the byte, and the receiving entity MUST restore the byte from its cached PHSF. If the bit value is a "0", the sending entity MUST NOT suppress the byte, and the receiving entity MUST restore the byte by using the next byte in the packet.

If this TLV is not included, the default is to suppress all bytes.

C.2.2.10.4 Payload Header Suppression Size (PHSS)

The value of this field is the total number of bytes in the Payload Header Suppression Field (PHSF) for a Service Flow that uses Payload Header Suppression.

Type	Length	Value
26.10	1	number of bytes in the suppression string

This TLV is used when a Service Flow is being created. For all packets that get classified and assigned to a Service Flow with Payload Header Suppression enabled, suppression MUST be performed over the specified number of bytes as indicated by the PHSS and according to the PHSM. If this TLV is included in a Service Flow definition with a value of 0 bytes, then Payload Header Suppression is disabled. A non-zero value indicates Payload Header Suppression is enabled. Until the PHSS value is known, the PHS rule is considered partially defined, and suppression will not be performed. A PHS rule becomes fully defined when both PHSS and PHSF are known.

C.2.2.10.5 Payload Header Suppression Verification (PHSV)

The value of this field indicates to the sending entity whether or not the packet header contents are to be verified prior to performing suppression. If PHSV is enabled, the sender MUST compare the bytes in the packet header with the bytes in the PHSF that are to be suppressed as indicated by the PHSM.

Type	Length	Value
26.11	1	0 = verify 1 = don't verify

If this TLV is not included, the default is to verify. Only the sender MUST verify suppressed bytes. If verification fails, the Payload Header MUST NOT be suppressed. (Refer to Section 8.4.3)

C.2.2.10.6 Vendor Specific PHS Parameters

This allows vendors to encode vendor-specific PHS parameters. The Vendor ID MUST be the first TLV embedded inside Vendor Specific PHS Parameters. If the first TLV inside Vendor Specific PHS Parameters is not a Vendor ID, then the TLV MUST be discarded. (Refer to Appendix C.1.1.17)

Type	Length	Value
26.43	n	

C.3 Encodings for Other Interfaces

C.3.1 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.3.2 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.4 Confirmation Code

The Confirmation Code (CC) provides a common way to indicate failures for Registration Response, Registration Ack, Dynamic Service Addition-Response, Dynamic Service Addition-Ack, Dynamic Service Delete-Response, Dynamic Service Change-Response, Dynamic Service Change-Ack and Dynamic Channel Change-Response MAC Management Messages. The confirmation codes in this section are used both as message Confirmation Codes and as Error Codes in Error Set Encodings which may be carried in these messages.

Confirmation Code is one of the following:

okay / success(0)
reject-other(1)
reject-unrecognized-configuration-setting(2)
reject-temporary / reject-resource(3)
reject-permanent / reject-admin(4)
reject-not-owner(5)
reject-service-flow-not-found(6)
reject-service-flow-exists(7)
reject-required-parameter-not-present(8)
reject-header-suppression(9)
reject-unknown-transaction-id(10)
reject-authentication-failure (11)
reject-add-aborted(12)
reject-multiple-errors(13)
reject-classifier-not-found(14)
reject-classifier-exists(15)
reject-PHS-rule-not-found(16)
reject-PHS-rule-exists(17)
reject-duplicate-reference-ID-or-index-in-message(18)
reject-multiple-upstream-service-flows(19)
reject-multiple-downstream-service-flows(20)
reject-classifier-for-another-service-flow(21)
reject-PHS-for-another-service-flow(22)
reject-parameter-invalid-for-context(23)
reject-authorization-failure(24)
reject-temporary-DCC(25)

The Confirmation Codes MUST be used in the following way:

- Okay or success(0) means the message was received and successful.
- Reject-other(1) is used when none of the other reason codes apply.
- Reject-unrecognized-configuration setting(2) is used when a configuration setting is not recognized or when its value is outside of the specified range.
- Reject-temporary(3), also known as reject-resource, indicates that the current loading of the CMTS or CM prevents granting the request, but that the request might succeed at another time.
- Reject-permanent(4), also known as reject-admin, indicates that, for policy, configuration, or capabilities reasons, the request would never be granted unless the CMTS or CM were manually reconfigured or replaced.
- Reject-not-owner(5) the requester is not associated with this service flow.
- Reject-service-flow-not-found(6) the Service Flow indicated in the request does not exist.
- Reject-service-flow-exists(7) the Service Flow to be added already exists.
- Reject-required-parameter-not-present(8) a required parameter has been omitted.

- Reject-header-suppression(9) the requested header suppression cannot be supported for whatever reason.
- Reject-unknown-transaction-id(10) the requested transaction continuation is invalid because the receiving end-point does not view the transaction as being 'in process' (*i.e.*, the message is unexpected or out of order).
- Reject-authentication-failure(11) the requested transaction was rejected because the message contained an invalid HMAC-digest, CMTS-MIC, provisioned IP address, or timestamp.
- Reject-add-aborted(12) the addition of a dynamic service flow was aborted by the initiator of the Dynamic Service Addition.
- Reject-multiple-errors (13) is used when multiple errors have been detected.
- Reject-classifier-not-found (14) is used when the request contains an unrecognized classifier ID.
- Reject-classifier-exists (15) indicates that the ID of a classifier to be added already exists.
- Reject-PHS-rule-not-found (16) indicates that the request contains an SFID/classifier ID pair for which no PHS rule exists.
- Reject-PHS-rule-exists (17) indicates that the request to add a PHS rule contains an SFID/classifier ID pair for which a PHS rule already exists.
- Reject-duplicate-reference-ID-or-index-in-message (18) indicates that the request used an SFR, classifier reference, SFID, or classifier ID twice in an illegal way.
- Reject-multiple-upstream-service-flows (19) is used when DSA/DSC contains parameters for more than one upstream flow.
- Reject-multiple-downstream-service-flows (20) is used when DSA/DSC contains parameters for more than one downstream flow.
- Reject-classifier-for-another-service-flow (21) is used in DSA-RSP when the DSA-REQ includes classifier parameters for a SF other than the new SF(s) being added by the DSA.
- Reject-PHS-for-another-service-flow (22) is used in DSA-RSP when the DSA-REQ includes a PHS rule for a SF other than the new SF(s) being added by the DSA.
- Reject-parameter-invalid-for-context(23) indicates that the parameter supplied cannot be used in the encoding in which it was included, or that the value of a parameter is invalid for the encoding in which it was included.
- Reject-authorization-failure(24) the requested transaction was rejected by the authorization module.
- Reject-temporary-DCC(25) indicates that the requested resources are not available on the current channels at this time, and the CM should re-request them on new channels after completing a channel change in response to a DCC command which the CMTS will send. If no DCC is received, the CM must wait for a time of at least T14 before re-requesting the resources on the current channels.

C.4.1 Confirmation Codes for Dynamic Channel Change

The CM may return in the DCC-RSP message an appropriate rejection code from Appendix C.4. It may also return one of the following Confirmation Codes which are unique to DCC-RSP.

depart(180)
arrive(181)
reject-already-there(182)

The Confirmation Codes MUST be used in the following way:

- depart(180) indicates the CM is on the old channel and is about to perform the jump to the new channel.
- arrive(181) indicates the CM has performed the jump and has arrived at the new channel.
- reject-already-there(182) indicates that the CMTS has asked the CM to move to a channel that it is already occupying.

C.4.2 Confirmation Codes for Major Errors

These confirmation codes MUST be used only as message Confirmation Codes in REG-ACK, DSA-RSP, DSA-ACK, DSC-RSP, or DSC-ACK messages, or as the Response code in REG-RSP messages for 1.1 CMs. In general, the errors associated with these confirmation codes make it impossible either to generate an error set that can be uniquely associated with a parameter set in the REG-REQ, DSA-REQ, or DSC-REQ message, or to generate a full RSP message.

reject-major-service-flow-error(200)
reject-major-classifier-error(201)
reject-major-PHS-rule-error(202)
reject-multiple-major-errors(203)
reject-message-syntax-error(204)
reject-primary-service-flow-error(205)
reject-message-too-big(206)
reject-invalid-modem-capabilities(207)

The Confirmation Codes MUST be used only in the following way:

- Reject-major-service-flow-error(200) indicates that the REQ message did not have either a SFR or SFID in a service flow encoding, and that service flow major errors were the only major errors.
- Reject-major-classifier-error(201) indicates that the REQ message did not have a classifier reference, or did not have both a classifier ID and a Service Flow ID, and that classifier major errors were the only major errors.
- Reject-major-PHS-rule-error(202) indicates that the REQ message did not have a both a Service Flow Reference/Identifier and a Classifier Reference/Identifier, and that PHS rule major errors were the only major errors.
- Reject-multiple-major-errors(203) indicates that the REQ message contained multiple major errors of types 200, 201, 202.
- Reject-message-syntax-error(204) indicates that the REQ message contained syntax error(s) (*e.g.*, a TLV length error) resulting in parsing failure.
- Reject-primary-service-flow-error(205) indicates that a REG-REQ or REG-RSP message did not define a required primary Service Flow, or a required primary Service Flow was not specified active.
- Reject-message-too-big(206) is used when the length of the message needed to respond exceeds the maximum allowed message size.
- Reject-invalid-modem-capabilities(207) indicates that the REG-REQ contained either that in invalid combination of modem capabilities or modem capabilities that are inconsistent with the services in the REG-REQ.

This page intentionally left blank.

Appendix D CM Configuration Interface Specification

D.1 CM IP Addressing

D.1.1 DHCP Fields Used by the 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) — to allow for the differentiation between DOCSIS 1.1 and DOCSIS 1.0 CM requests, a compliant CM MUST send the following ASCII coded string in Option code 60, "docsis1.1:xxxxxxx". Where xxxxx MUST be an ASCII representation of the hexadecimal encoding of the Modem Capabilities, refer to C.1.3.1. For example, the ASCII encoding for the first two TLVs (concatenation and DOCSIS Version) of a DOCSIS 1.1 modem would be 05nn010101020101. Note that many more TLVs are required for a DOCSIS1.1 modem and the field "nn" will contain the length of all the TLVs. This example shows only two TLVs for simplicity.
- 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. Fields identified as critical MUST be present in the DHCP response, and fields identified as non-critical SHOULD be present. The CM MUST configure itself with the critical fields from the DHCP response, and, if present, with the non-critical fields.

- The IP address to be used by the CM (yiaddr) (critical).
- The IP address of the TFTP server for use in the next phase of the bootstrap process (siaddr) (critical).
- 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 (non-critical).
- The name of the CM configuration file to be read from the TFTP server by the CM (file) (critical).
- The subnet mask to be used by the CM (Subnet Mask, option 1) (non-critical).
- 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 (non-critical).

- 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, but MUST use at least one (non-critical).
- A list of [RFC-868] time-servers from which the current time may be obtained (Time Server Option, option 4) (non-critical).
- A list of SYSLOG servers to which logging information may be sent (Log Server Option, option 7); see [DOCSIS5] (non-critical).

If a critical field is missing or is invalid in the DHCP response during initialisation, the CM MUST log an error and reinitialise its MAC and continue channel scanning.

If a non-critical field is missing or is invalid in the DHCP response during initialisation, the CM MUST log a warning, ignore the field and go operational, with the following considerations:

- If the subnet mask is missing or is invalid, the CM MUST use the default for the IP of Class A, B or C as defined in [RFC-791].
- If the time server is missing or is invalid, the CM MUST initialise the time for the events to Jan 1, 1970, 0h00.

If the IP address field is missing or is invalid in the DHCP response during renew or rebind, the CM MUST log an error and reinitialise its MAC and continue channel scanning.

If any other critical or non-critical field is missing or is invalid in the DHCP response during renew or rebind, the CM MUST log a warning, ignore the field if it is invalid and stay operational.

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. Bridging CMTSs SHOULD also provide this functionality.

D.2 CM Configuration

D.2.1 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. To allow uniform management of CM's conformant to this specification, conformant CM's MUST support a 8192-byte configuration file at a minimum.

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 D-1:

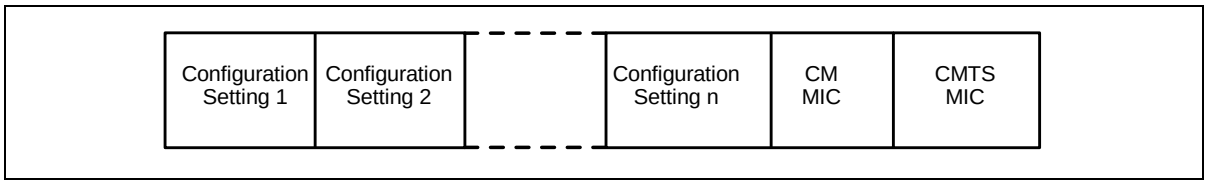


Figure D-1. Binary Configuration File Format

D.2.2 Configuration File Settings

The following configuration settings MUST be included in the configuration file and MUST be supported by all CMs. The CM MUST NOT send a REG-REQ based on a configuration file that lacks these mandatory items.

- Network Access Configuration Setting
- CM MIC Configuration Setting
- CMTS MIC Configuration Setting
- End Configuration Setting
- DOCSIS 1.0 Class of Service Configuration Setting

— or —

- Upstream Service Flow Configuration Setting
- Downstream Service Flow Configuration Setting

Note: A DOCSIS 1.0 CM must be provided with a DOCSIS 1.0 Class of Service Configuration. A CM conformant with this specification should only be provisioned with DOCSIS 1.0 Class of Service Configuration information if it is to behave as a DOCSIS 1.0 CM, otherwise it should be provisioned with Service Flow Configuration Settings

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
- Upstream Packet Classification Setting
- Downstream Packet Classification Setting
- SNMP Write-Access Control
- SNMP MIB Object
- Software Server IP Address
- CPE Ethernet MAC Address
- Maximum Number of CPEs
- Maximum Number of Classifiers
- Privacy Enable Configuration Setting
- Payload Header Suppression
- TFTP Server Timestamp
- TFTP Server Provisioned Modem Address
- Pad Configuration Setting
- SNMPv3 Notification Receiver

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

Note: There is a limit on the size of registration request and registration response frames (see Section 6.2.5.2). The configuration file should not be so large as to require the CM or CMTS to exceed that limit.

D.2.3 Configuration File Creation

The sequence of operations required to create the configuration file is as shown in Figure D-2 through Figure D-5.

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

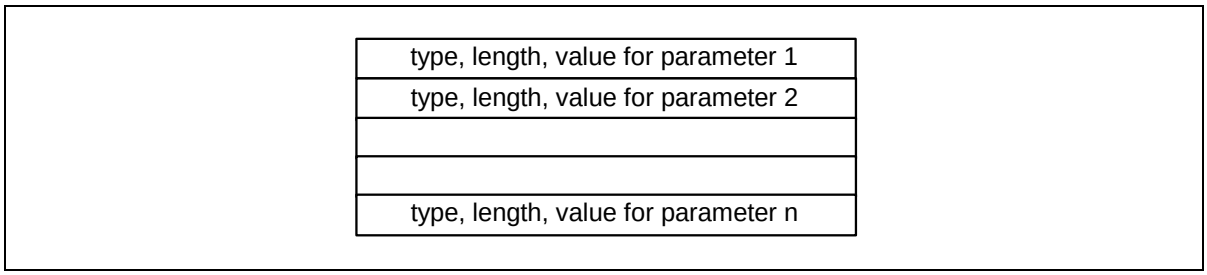


Figure D-2. Create TLV Entries for Parameters Required by the CM

2. Calculate the CM message integrity check (MIC) configuration setting as defined in Section D.2.3.1 and add to the file following the last parameter using code and length values defined for this field.

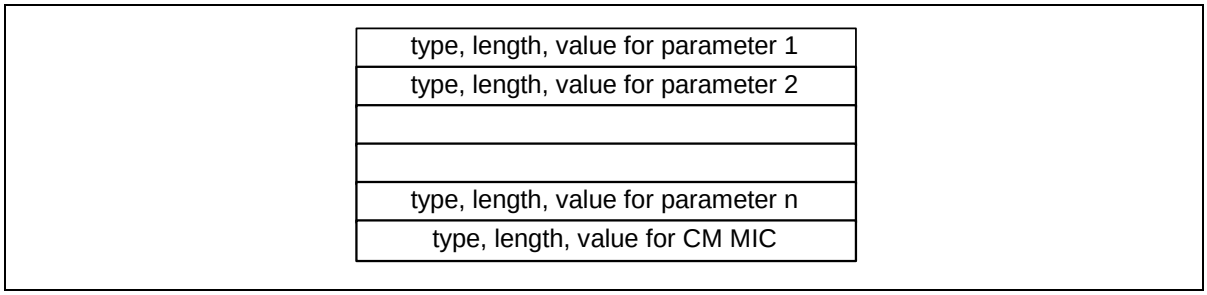


Figure D-3. Add CM MIC

3. Calculate the CMTS message integrity check (MIC) configuration setting as defined in Section D.3.1 and add to the file following the CM MIC using code and length values defined for this field.

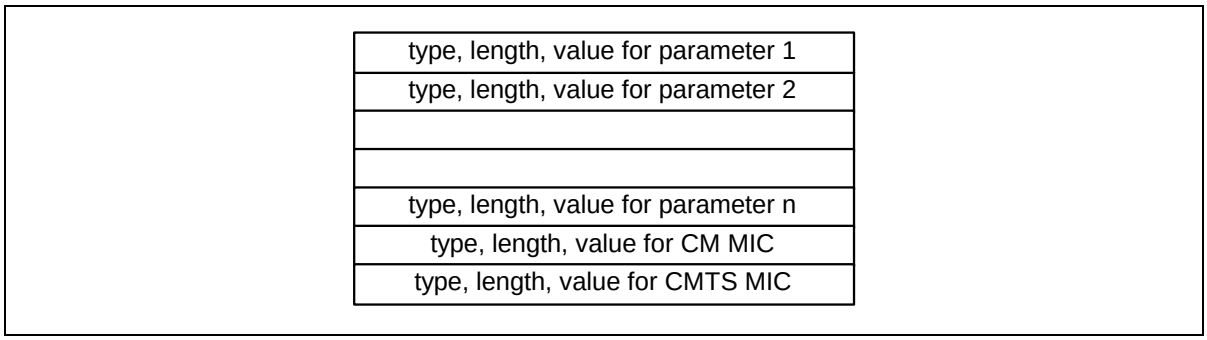


Figure D-4. Add CMTS MIC

4. Add the end of data marker.

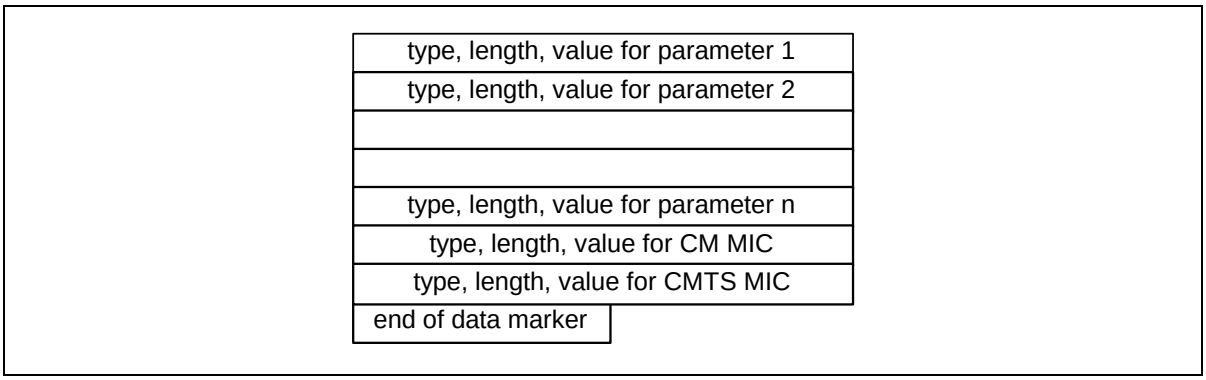


Figure D-5. Add End of Data Marker

D.2.3.1 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 then the configuration file MUST be discarded

D.3 Configuration Verification

It is necessary to verify that the CM's configuration file has come from a trusted source. Thus, the CMTS and the configuration server share an Authentication String that they use to verify portions of the CM's configuration in the Registration Request.

D.3.1 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
- DOCSIS 1.0 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
- Upstream Packet Classification Setting
- Downstream Packet Classification Setting
- Upstream Service Flow Configuration Setting
- Downstream Service Flow Configuration Setting
- Maximum Number of Classifiers
- Privacy Enable Configuration Setting
- Payload Header Suppression
- Subscriber Management Control
- Subscriber Management CPE IP Table
- Subscriber Management Filter Groups

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 Appendix D.3.1.1.

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.

D.3.1.1 Digest Calculation

The CMTS MIC digest field MUST be calculated using HMAC-MD5 as defined in [RFC-2104].

This page intentionally left blank.

Appendix E MAC Service Definition

This section is informational. In case of conflict between this section and any normative section of this specification, the normative section takes precedence.

E.1 MAC Service Overview

The DOCSIS MAC provides a protocol service interface to upper-layer services. Examples of upper-layer services include a DOCSIS bridge, embedded applications (*e.g.*, Packetcable/VOIP), a host interface (*e.g.*, NIC adapter with NDIS driver), and layer three routers (*e.g.*, IP router).

The MAC Service interface defines the functional layering between the upper layer service and the MAC. As such it defines the functionality of the MAC which is provided by the underlying MAC protocols. This interface is a protocol interface, not a specific implementation interface.

The following data services are provided by the MAC service interface:

- A MAC service exists for classifying and transmitting packets to MAC service flows.
- A MAC service exists for receiving packets from MAC service flows. Packets may be received with suppressed headers.
- A MAC service exists for transmitting and receiving packets with suppressed headers. The headers of transmitted packets are suppressed based upon matching classifier rules. The headers of received suppressed packets are regenerated based upon a packet header index negotiated between the CM and CMTS.
- A MAC service exists for synchronization of grant timing between the MAC and the upper layer service. This clock synchronization is required for applications such as embedded Packetcable VOIP clients in which the packetization period needs to be synchronized with the arrival of scheduled grants from the CMTS.
- A MAC service exists for synchronization of the upper layer clock with the CMTS Controlled Master Clock.

It should be noted that a firewall and policy based filtering service may be inserted between the MAC layer and the upper layer service, but such a service is not modeled in this MAC service definition.

The following control services are provided by the MAC service interface:

- A MAC service exists for the upper layer to learn of the existence of provisioned service flows and QoS traffic parameter settings at registration time.
- A MAC service exists for the upper layer to create service flows. Using this service the upper layer initiates the admitted/activated QoS parameter sets, classifier rules, and packet suppression headers for the service flow.
- A MAC service exists for the upper layer to delete service flows.
- A MAC service exists for the upper layer to change service flows. Using this service the upper layer modifies the admitted/activated QoS parameter sets, classifier rules, and packet suppression headers.
- A MAC service exists for controlling the classification of and transmission of PDUs with suppressed headers. At most a single suppressed header is defined for a single classification rule. The upper layer service is responsible for defining both the definition of suppressed headers (including wild-card don't-suppress fields) and the unique classification rule that discriminates each header. In addition to the classification rule, the MAC service can perform a full match of all remaining header bytes to prevent generation of false headers if so configured by the upper layer service.
- A MAC service exists for controlling two-phase control of QoS traffic resources. Two phase activation is controlled by the upper layer service provide both admitted QoS parameters and active QoS parameters within the appropriate service request. Upon receipt of an affirmative indication the upper layer service

knows that the admitted QoS parameter set has been reserved by the CMTS, and that the activated QoS parameter set has been activated by the CMTS. Barring catastrophic failure (such as resizing of the bandwidth of the upstream PHY), admitted resources will be guaranteed to be available for activation, and active resources will be guaranteed to be available for use in packet transmission.

A control function for locating an unused service flow and binding it or a specific identified service flow to a specific upper layer service may also exist. The details of such a function are not specified and are implementation dependent.

Other control functions may exist at the MAC service interface, such as functions for querying the status of active service flows and packet classification tables, or functions from the MAC service to the upper layer service to enable the upper layer service to authorize service flows requested by the peer MAC layer service, but those functions are not modeled in this MAC service definition.

Other MAC services that are not service flow related also exist, such as functions for controlling the MAC service MAC address and SAID multicast filtering functions, but those functions are not modeled in this MAC service definition.

E.1.1 MAC Service Parameters

The MAC service utilizes the following parameters. For a full description of the parameters consult the Theory of Operation and other relevant sections within the body of the RFI specification.

- Service Flow QoS Traffic Parameters

MAC activate-service-flow and change-service-flow primitives allow common, upstream, and downstream QoS traffic parameters to be provided. When such parameters are provided they override whatever values were configured for those parameters at provisioning time or at the time the service flow was created by the upper layer service.

- Active/Admitted QoS Traffic Parameters

If two-phase service flow activation is being used, then two complete sets of QoS Traffic Parameters are controlled. The admitted QoS Parameters state the requirements for reservation of resources to be authorized by the CMTS. The activated QoS Parameters state the requirements for activation of resources to be authorized by the CMTS. Admitted QoS parameters may be activated at a future time by the upper layer service. Activated QoS parameters may be used immediately by the upper layer service.

- Service Flow Classification Filter Rules

Zero or more classification filter rules may be provided for each service flow that is controlled by the upper layer service. Classifiers are identified with a classifier identifier.

- Service Flow PHS Suppressed Headers

Zero or more PHS suppressed header strings with their associated verification control and mask variables may be defined for each service flow. When such headers are defined, they are associated 1-to-1 with specific classification rules. In order to regenerate packets with suppressed headers a payload header suppression index is negotiated between the CM and CMTS.

E.2 MAC Data Service Interface

MAC services are defined for transmission and reception of data to and from service flows. Typically an upper layer service will utilize service flows for mapping of various classes of traffic to different service flows. Mappings to service flows may be defined for low priority traffic, high priority traffic, and multiple special traffic classes such as constant bit rate traffic which is scheduled by periodic grants from the CMTS at the MAC layer.

The following specific data service interfaces are provided by the MAC service to the upper layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

MAC_DATA.request
MAC_DATA.indicate
MAC_GRANT_SYNCHRONIZE.indicate
MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate

E.2.1 MAC_DATA.request

Issued by the upper-layer service to request classification and transmission of an IEEE 802.3 or DIX formatted PDU to the RF.

Parameters:

- PDU - IEEE 802.3 or DIX encoded PDU including all layer two header fields and optional FCS. PDU is the only mandatory parameter.
- padding - is used when the PDU is less than 60 bytes and it is desired to maintain [ISO8802-3] transparency.
- ServiceFlowID - if included the MAC service circumvents the packet classification function and maps the packet to the specific service flow indicated by the ServiceFlowID value.
- ServiceClassName, RulePriority - if included this tuple identifies the service class name of an active service flow to which the packet is to be mapped so long as a classifier does not exist at a rule priority higher than the rule priority supplied.

Expanded Service Description:

Transmit a PDU from upper-layer service to MAC/PHY. The only mandatory parameter is PDU. PDU contains all layer-2 headers, layer-3 headers, data, and (optional) layer-2 checksum.

If PDU is the only parameter, the packet is subjected to the MAC packet classification filtering function in order to determine how the packet is mapped to a specific service flow. The results of the packet classification operation determine on which service flow the packet is to be transmitted and whether or not the packet should be transmitted with suppressed headers.

If the parameter ServiceFlowID is supplied the packet can be directed to the specifically identified service flow.

If the parameter tuple ServiceClassName, RulePriority is supplied the packet is directed to the first active service flow that matches the service class name so long as a classifier does not exist at a rule priority higher than the rule priority supplied. This service is used by upper layer policy enforcers to allow zero or more dynamic rules to be matched for selected traffic (*e.g.*, voice) while all other traffic is forced to a service flow within the named ServiceFlowClass. If no active service flow with the Service Class Name exists, then the service perform normal packet classification.

In all cases, if no classifier match is found, or if none of the combinations of parameters maps to a specific service flow, the packet will be directed to the primary service flow.

The following pseudo code describes the intended operation of the MAC_DATA.request service interface:

```
MAC_DATA.request
  PDU
  [ServiceFlowID]
  [ServiceClassName, RulePriority]

FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName) returns ServiceFlowID of first service flow whose
ServiceClassName equals the parameter of the procedure or NULL if no matching service flow found.

SEARCH_CLASSIFIER_TABLE (PriorityRange) searches all rules within the specified priority range and
returns either the ServiceFlowID associated with the rule or NULL if no classifier rule found.

TxServiceFlowID = NULL

IF  (ServiceFlowID DEFINED)
          TxServiceFlowID = MAC_DATA.ServiceFlowID

ELSEIF (ServiceClassName DEFINED and RulePriority DEFINED)
          TxServiceFlowID = FIND_FIRST_SERVICE_FLOW_ID (ServiceClassName)
          SearchID = SEARCH_CLASSIFIER_TABLE (All Priority Levels)
          IF (SearchID not NULL and ClassifierRule.Priority >= MAC_DATA.RulePriority)
                    TxServiceFlowID = SearchID

ELSE  [PDU only]
         TxServiceFlow = SEARCH_CLASSIFIER_TABLE (All Priority Levels)

IF (TxServiceFlowID = NULL)
     TRANSMIT_PDU (PrimaryServiceFlowID)
ELSE
      TRANSMIT_PDU (TxServiceFlowID)
```

E.2.2 MAC_DATA.indicate

Issued by the MAC to indicate reception of an IEEE 802.3 or DIX PDU for the upper-layer service from the RF.

Parameters:

- PDU - IEEE 802.3 or DIX encoded PDU including all layer two header fields and FCS.

E.2.3 MAC_GRANT_SYNCHRONIZE.indicate

Issued by the MAC service to the upper layer service to indicate the timing of grant arrivals from the CTMS. It is not stated how the upper layer derives the latency if any between the reception of the indication and the actual arrival of grants (within the bounds of permitted grant jitter) from the CMTS. It should be noted that in UGS applications it is expected that the MAC layer service will increase the grant rate or decrease the grant rate based upon the number of grants per interval QoS traffic parameter. It should also be noted that as the number of grants per interval is increased or decreased that the timing of grant arrivals will change also. It should also be noted that when synchronization is achieved with the CMTS downstream master clock, this indication may only be required once per active service flow. No implication is given as to how this function is implemented.

Parameters:

- ServiceFlowID - unique identifier value for the specific active service flow receiving grants.

E.2.4 MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate

Issued by the MAC service to the upper layer service to indicate the timing of the CMTS master clock. No implication is given as to how often or how many times this indication is delivered by the MAC service to the upper layer service. No implication is given as to how this function is implemented.

Parameters:

- No parameters specified.

E.3 MAC Control Service Interface

A collection of MAC services are defined for control of MAC service flows and classifiers. It should be noted that an upper layer service may use these services to provide an upper layer traffic construct such as "connections" or "subflows" or "micro-flows". However, except for the ability to modify individual classifiers, no explicit semantics is defined for such upper layer models. Thus control of MAC service flow QoS parameters is specified in the aggregate.

The following specific control service interface functions are provided by the MAC service to the upper layer service. These represent an abstraction of the service provided and do not imply a particular implementation:

MAC_REGISTRATION_RESPONSE.indicate
MAC_CREATE_SERVICE_FLOW.request/response/indicate
MAC_DELETE_SERVICE_FLOW.request/response/indicate
MAC_CHANGE_SERVICE_FLOW.request/response/indicate

E.3.1 MAC_REGISTRATION_RESPONSE.indicate

Issued by the DOSCIS MAC to the upper layer service to indicate the complete set service flows and service flow QoS traffic parameters that have been provisioned and authorized by the registration phase of the MAC. Subsequent changes to service flow activation state or addition and deletion of service flows are communicated to the upper layer service with indications from the other MAC control services.

Parameters:

- Registration TLVs - any and all TLVs that are needed for service flow and service flow parameter definition including provisioned QoS parameters. See the normative body of the specification for more details.

E.3.2 MAC_CREATE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the creation of a new service flow within the MAC service. This primitive is not issued for service flows that are configured and registered, but rather for dynamically created service flows. This primitive may also define classifiers for the service flow and supply admitted and activated QoS parameters. This function invokes DSA signaling.

Parameters:

- ServiceFlowID - unique id value for the specific service flow being created.
- ServiceClassName - service flow class name for the service flow being created.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.

- Service Flow Payload Header Suppression Rules - Zero or more PHS rules for each service flow that is controlled by the upper layer service.
- Service Flow Classification Filter Rules - Zero or more classification filter rules for each service flow that is controlled by the upper layer service. Classifiers are identified with a classifier identifier.

E.3.3 MAC_CREATE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to create a service flow.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being created.
- ResponseCode - success or failure code

E.3.4 MAC_CREATE_SERVICE_FLOW.indicate

Issued by the MAC service to notify the upper-layer service of the creation of a new service flow within the MAC service. This primitive is not issued for service flows that have been administratively pre-configured, but rather for dynamically defined service flows. In the present document, this notification is advisory only.

Parameters:

- ServiceFlowID - unique id value for the specific service flow being created.
- ServiceClassName - service flow class name for the service flow being created.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Service Flow Payload Header Suppression Rules - Zero or more PHS rules for each service flow that is controlled by the upper layer service.
- Service Flow Classification Filter Rules - Zero or more classification filter rules for each service flow that is controlled by the upper layer service. Classifiers are identified with a classifier identifier.

E.3.5 MAC_DELETE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request the deletion of a service flow and all QoS parameters including all associated classifiers and PHS rules. This function invokes DSD signaling.

Parameters:

- ServiceFlowID(s) - unique identifier value(s) for the deleted service flow(s).

E.3.6 MAC_DELETE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to delete a service flow.

Parameters:

- ResponseCode - success or failure code

E.3.7 MAC_DELETE_SERVICE_FLOW.indicate

Issued by the MAC service to notify the upper-layer service of deletion of a service flow within the MAC service.

Parameters:

- ServiceFlowID(s) - unique identifier value(s) for the deleted service flow(s).

E.3.8 MAC_CHANGE_SERVICE_FLOW.request

Issued by the upper-layer service to the MAC to request modifications to a specific created and acquired service flow. This function is able to define both the complete set of classifiers and incremental changes to classifiers (add/remove). This function defines the complete set of admitted and active QoS parameters for a service flow. This function invokes DSC MAC-layer signaling.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being modified.
- zero or more packet classification rules with add/remove semantics and LLC, IP, and 802.1pq parameters.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Service Flow Payload Header Suppression Rules - Zero or more PHS rules for each service flow that is controlled by the upper layer service.

E.3.9 MAC_CHANGE_SERVICE_FLOW.response

Issued by the MAC service to the upper layer service to indicate the success or failure of the request to change a service flow.

Parameters:

- ServiceFlowID - unique identifier value for the specific service flow being released.
- ResponseCode - success or failure code

E.3.10 MAC_CHANGE_SERVICE_FLOW.indicate

Issued by the DOSCIS MAC service to notify upper-layer service of a request to change a service flow. In this specification the notification is advisory only and no confirmation is required before the service flow is changed. Change-service-flow indications are generated based upon DSC signaling. DSC signaling can be originated based upon change-service-flow events between the peer upper-layer service and its MAC service, or based upon network resource failures such as a resizing of the total available bandwidth at the PHY layer. How the upper layer service reacts to forced reductions in admitted or reserved QoS traffic parameters is not specified.

Parameters:

- ServiceFlowID - unique identifier for the service flow being activated.
- packet classification rules with LLC, IP, and 802.1pq parameters, and with zero or more PHS_CLASSIFIER_IDENTIFIERs.
- Admitted QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Activated QoS Parameters - zero or more upstream, downstream, and common traffic parameters for the service flow.
- Service Flow Payload Header Suppression Rules - Zero or more PHS rules for each service flow that is controlled by the upper layer service.

E.4 MAC Service Usage Scenarios

Upper layer entities utilize the services provided by the MAC in order to control service flows and in order to send and receive data packets. The partition of function between the upper-layer-service and the MAC service service is demonstrated by the following scenarios.

E.4.1 Transmission of PDUs from Upper Layer Service to MAC DATA Service

- Upper layer service transmits PDUs via the MAC_DATA service.
- MAC_DATA service classifies transmitted PDUs using the classification table, and transmits the PDUs on the appropriate service flow. The classification function may also cause the packet header to be suppressed according to a header suppression template stored with the classification rule. It is possible for the upper layer service to circumvent this classification function.
- MAC_DATA service enforces all service flow based QoS traffic shaping parameters.
- MAC_DATA service transmits PDUs on DOCSIS RF as scheduled by the MAC layer.

E.4.2 Reception of PDUs to Upper Layer Service from MAC DATA Service

- PDUs are received from the DOCSIS RF.
- If PDU is sent with a suppressed header, the header is regenerated before the packet is subjected to further processing.
- In the CMTS the MAC_DATA service classifies PDUs ingress from the RF using the classification table and then polices the QoS traffic shaping and validates addressing as performed by the CM. In the CM no per-packet service flow classification is required for traffic ingress from the RF.
- Upper layer service receives PDUs from the MAC_DATA.indicate service.

E.4.3 Sample Sequence of MAC Control and MAC Data Services

A possible CM-oriented sequence of MAC service functions for creating, acquiring, modifying, and then using a specific service flow is as follows:

- MAC_REGISTER_RESPONSE.indicate

  Learn of any provisioned service flows and their provisioned QoS traffic parameters.

- MAC_CREATE_SERVICE_FLOW.request/response

  Create new service flow. This service interface is utilized if the service flow was learned as not provisioned by the MAC_REGISTER_RESPONSE service interface. Creation of a service flow invokes DSA signaling.

- MAC_CHANGE_SERVICE_FLOW.request/response

  Define admitted and activated QoS parameter sets, classifiers, and packet suppression headers. Change of a service flow invokes DSC signaling.

- MAC_DATA.request

  Send PDUs to MAC service for classification and transmission.

- MAC_DATA.indication

  Receive PDUs from MAC service.

- MAC_DELETE_SERVICE_FLOW.request/response

  Delete service flow. Would likely be invoked only for dynamically created service flows, not provisioned service flows. Deletion of a service flow uses DSD signaling.

This page intentionally left blank.

Appendix F Example Preamble Sequence

F.1 Introduction

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 6.3.3, Table 6-18), but with each burst profile able to specify the start location within this sequence of bits and the length of the preamble (Section 6.3.3, Table 6-19). 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 first bit of the Preamble Pattern is the first bit into the symbol mapper (Figure 4-9), and is I1 in the first symbol of the burst (see Section 4.2.2.2). As an example, per Table 6-19, 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 F.2.

F.2 Example Preamble Sequence

The following is the example 1024-bit preamble sequence:

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

Appendix G DOCSIS v1.0/v1.1 Interoperability

G.1 Introduction

This specification is informally referred to as DOCSIS 1.1. It is the second generation of DOCSIS 1.0 specified in [SCTE1]. The terms DOCSIS 1.1 and DOCSIS 1.0 refer to these two different specifications.

The DOCSIS 1.1 specification, primarily aims at enhancing the limited QoS functionality of a DOCSIS 1.0 based cable access system. New MAC messages have been defined for dynamic QoS signaling, and several new QoS parameter encodings have been defined in the existing MAC messages. A DOCSIS 1.1 CMTS can better support the requirements of delay/jitter sensitive traffic on a DOCSIS 1.1 CM.

Besides supporting a rich set of QoS features for DOCSIS 1.1 CMs, the DOCSIS 1.1 CMTS must be backwards compatible with a DOCSIS 1.0 CM. Furthermore, it is necessary for a 1.1 CM to function like a 1.0 CM when interoperating with a 1.0 CMTS.

This section describes the interoperability issues and trade-offs involved, when the operator wishes to support DOCSIS 1.0 as well as DOCSIS 1.1 CMs on the same cable access channel.

G.2 General Interoperability Issues

This section addresses the general DOCSIS 1.0/DOCSIS 1.1 interoperability issues that do not impact the performance during normal operation of the CMs.

G.2.1 Provisioning

The parameters of the TFTP config file for a DOCSIS 1.1 CM, are a superset of those for a DOCSIS 1.0 CM. Configuration file editors will have to be enhanced to incorporate support for these new parameters and the new MIC calculation.

A TFTP configuration file containing DOCSIS 1.0 Class of Service TLVs is considered a “DOCSIS 1.0-style” configuration file. A TFTP configuration file containing DOCSIS 1.1 Service Flow TLVs is considered a “DOCSIS 1.1-style” configuration file. A TFTP configuration file containing both Class of Service and Service Flow TLVs will be rejected by the CMTS (see Section 9.2.9).

If a DOCSIS 1.1 CM is provisioned with a DOCSIS 1.0-style TFTP configuration file, like a DOCSIS 1.0 CM, it MUST NOT respond to REG-RSP with REG-ACK (although in the REG-REQ it MUST still specify “DOCSIS v1.1” in the DOCSIS Version Modem Capability and MAY specify 1.1 Modem Capabilities that it can support when registered like a DOCSIS 1.0 CM). Thus, a DOCSIS 1.1 CM can be provisioned to work seamlessly on either a DOCSIS 1.0 or a DOCSIS 1.1 network.

If a DOCSIS 1.1 CM supports certain 1.1 capabilities when registered like a DOCSIS 1.0 CM (as indicated by the Modem Capabilities Encoding), those features MUST function according to the requirements defined in the DOCSIS 1.1 specifications.

On the other hand, DOCSIS 1.0 CMs do not recognize (and ignore) many of the new TLVs in a DOCSIS 1.1-style config file, and will be unable to register successfully if provisioned with a DOCSIS 1.1 configuration file. To prevent any functionality mismatches, a DOCSIS 1.1 CMTS MUST reject any Registration Request with DOCSIS 1.1-specific configuration parameters that are not supported by the associated Modem Capabilities encoding in the REG-REQ (see Appendix C.1.3.1).

G.2.2 Registration

A DOCSIS 1.1 CMTS is designed to handle the existing registration TLVs from DOCSIS 1.0 CMs as well as the new TLVs (namely, types 22 to 30) from the DOCSIS 1.1 CM.

There is a slight difference in the Registration related messaging procedure when the DOCSIS 1.1 CMTS is responding to a DOCSIS 1.1 CM as opposed to DOCSIS 1.0 CM. A DOCSIS 1.1 CM could be configured to use the Service Class Name which is statically defined at the CMTS instead of asking for the service class parameters explicitly. When such a Registration-Request is received by the DOCSIS 1.1 CMTS, it encodes the actual parameters of that service class in the Registration-Response and expects the DOCSIS 1.1 specific Registration-Acknowledge MAC message from the CM. If the detailed capabilities in the Registration-Response message exceed those the CM is capable of supporting, the CM is required to indicate this to the CMTS in its Registration-Acknowledge.

When a DOCSIS 1.0 CM registers with the same CMTS, the default DOCSIS 1.0 version is easily identified by the absence of the “DOCSIS Version” Modem Capabilities encoding in the Registration-Request. The Registration-Request from DOCSIS 1.0 CM explicitly requests all non-default service class parameters in the Registration-Request per its provisioning information. Absence of a Service Class Names eliminates the need for the DOCSIS 1.1 CMTS to explicitly specify the service class parameters in the Registration-Response using DOCSIS 1.1 TLVs. When a DOCSIS 1.1 CMTS receives a Registration-Request containing DOCSIS 1.0 Class of Service Encodings, it will respond with the regular DOCSIS 1.0-style Registration-Response and not expect the CM to send the Registration-Acknowledge MAC message.

Another minor issues is that a DOCSIS 1.0 CM will request for a bi-directional (with Upstream/Downstream parameters) service class from the CMTS using a Class-of-Service Configuration Setting.

Since DOCSIS 1.1 CMTS typically operates with unidirectional service classes, it can easily translate a DOCSIS 1.0 Class-of-Service Configuration Setting into DOCSIS 1.1 Service Flow Encodings for setting up unidirectional service classes in local QoS implementation. However, for DOCSIS 1.0 modems, the DOCSIS 1.1 CMTS MUST continue to maintain the QoSProfile table (with bi-directional Class parameters) for backward compatibility with DOCSIS 1.0 MIB.

Thus, if properly provisioned, a DOCSIS 1.0 and a DOCSIS 1.1 CM can successfully register with the same DOCSIS 1.1 CMTS. Likewise, a DOCSIS 1.0 and a DOCSIS 1.1 CM can successfully register with the same DOCSIS 1.0 CMTS.

G.2.3 Dynamic Service Establishment

There are 8 new MAC messages that relate to Dynamic Service Establishment. A DOCSIS 1.0 CM will never send them to any CMTS since they are unsupported. A DOCSIS 1.1 CM will never send them to a DOCSIS 1.0 CMTS because (a) to register successfully it has to be provisioned as a DOCSIS 1.0 CM and (b) when provisioned as a DOCSIS 1.0 CM it acts identically. When a DOCSIS 1.1 CM is connected to a DOCSIS 1.1 CMTS these messages work as expected.

G.2.4 Fragmentation

Fragmentation is initiated by the CMTS. Thus, a DOCSIS 1.0 CMTS will never initiate fragmentation since it knows nothing about it. A DOCSIS 1.1 CMTS can only initiate fragmentation for DOCSIS 1.1 CMs. A DOCSIS 1.1 CMTS MUST NOT attempt to fragment transmissions from a DOCSIS v1.0 CM that has not indicated a Modem Capabilities encoding for Fragmentation Support with a value of 1.

G.2.5 Multicast Support

It is mandatory for DOCSIS 1.0 CM's to support forwarding of multicast traffic. However, the specification is silent on IGMP support. Thus, the only standard mechanism for controlling IP-multicast on DOCSIS 1.0 CMs is through SNMP and packet filters. Designers of DOCSIS 1.0 networks will have to deal with these limitations and expect no different from DOCSIS 1.0 CM's on a DOCSIS 1.1 network.

G.2.6 Upstream Channel Change

A DOCSIS 1.1 CMTS is capable of specifying the level of re-ranging to be performed when it issues an UCC-Request to the CM. This re-ranging technique parameter is specified by the DOCSIS 1.1 CMTS using a new TLV in the UCC-Request MAC message.

DOCSIS 1.1 CMs that recognize this new TLV in the UCC-Request can benefit by only re-ranging to the level specified by this TLV. This can help in reducing the reinitialization time following a UCC, for the DOCSIS 1.1 CM carrying a voice call. A DOCSIS 1.1 CMTS is aware of the type of CM to which it is issuing the UCC-Request. It can refrain from inserting this re-ranging TLV in the UCC-Request for DOCSIS 1.0 CMs. If a DOCSIS 1.1 CMTS inserts this re-ranging TLV in the UCC-Request, the DOCSIS 1.0 CMs which do not recognize this TLV will ignore its contents and perform the default DOCSIS 1.0 re-ranging from start (Initial-Maintenance). The DOCSIS 1.1 CMTS accepts default initial ranging procedure from any modem issued the UCC-Request.

Thus DOCSIS 1.0 and DOCSIS 1.1 CMs on the same upstream channel can be individually requested to change upstream channels without any interoperability issues caused by the DOCSIS 1.1-style re-ranging TLV in the UCC-request.

G.3 Hybrid Devices

Some DOCSIS 1.0 CM designs may be capable of supporting individual DOCSIS 1.1 features via a software upgrade. Similarly, some DOCSIS 1.0 CMTS's may be capable of supporting individual DOCSIS 1.1 features. To facilitate these "hybrid" devices, the majority of DOCSIS 1.1 features are individually enumerated in the Modem Capabilities.

DOCSIS 1.0 hybrid CM's MAY request DOCSIS 1.1 features via this mechanism. However, unless a CM is fully DOCSIS 1.1 compliant (*i.e.,* not a hybrid), it MUST NOT send a "DOCSIS Version" Modem Capability which indicates anything besides DOCSIS 1.0.

If a hybrid CM intends to request such 1.1 capabilities from the CMTS during registration, it MUST send the ASCII coded string in Option code 60 of its DHCP request, "docsis1.0:xxxxxxx". Where xxxxx MUST be an ASCII representation of the hexadecimal encoding of the Modem Capabilities, refer to Appendix C.1.3.1 and Appendix D.1.1. The DHCP server MAY use such information to determine what configuration file the CM is to use.

In order to control the hybrid operation of modems, if a DOCSIS 1.1 CMTS receives a 1.0-style Registration Request message from a CM, the CMTS MUST, by default, force the modem to operate in a "pure" 1.0 mode with respect to certain features by disabling those features via the Modem Capabilities Encoding in the Registration Response. Specifically, the CMTS MUST support the six default values given in square brackets in Table G-1. The CMTS MAY provide switches, as indicated in Table G-1, for the operator to selectively allow certain hybrid features to be enabled.

Table G-1. Hybrid Mode Controls

	Concatenation Support	Fragmentation Support	Privacy Support
1.0 CM	allow/[deny]	allow/[deny]	allow BPI+/[force BPI]
1.1 CM in 1.0 mode	allow/[deny]	allow/[deny]	allow BPI+/[force BPI]

Normally, a DOCSIS 1.0 CMTS would set all unknown Modem Capabilities to 'Off' in the Registration Response indicating that these features are unsupported and MUST NOT be used by the CM. A DOCSIS 1.0 hybrid CMTS's MAY leave supported Modem Capabilities set to 'On' in the Registration Response. However, unless a CMTS is fully DOCSIS 1.1 compliant (*i.e.*, not a hybrid), it MUST still set all "DOCSIS Version" Modem Capabilities to DOCSIS 1.0.

As always, any Modem Capability set to 'Off' in the Registration Response must be viewed as unsupported by the CMTS and MUST NOT be used by the CM.

G.4 Interoperability & Performance

This section addresses the issue of performance impact on the QoS for DOCSIS 1.1 CMs when DOCSIS 1.0 and DOCSIS 1.1 CMs are provisioned to share the same upstream MAC channel.

The DOCSIS 1.0 CMs lack the ability to explicitly set their request policy (or provide scheduling parameters) for the advanced DOCSIS 1.1 scheduling mechanisms like "Unsolicited Grant Service" and "Real-Time Polling Service". Thus, DOCSIS 1.0 CMs will only receive statically configured "Tiered Best Effort" or "CIR" service on the upstream. The DOCSIS 1.1 CMs on the same upstream channel can explicitly request for additional Service Flows when required, using the DOCSIS 1.1 DSA-Request MAC message. Thus, DOCSIS 1.1 CMs can benefit from the advanced scheduling mechanisms of a DOCSIS 1.1 CMTS for their real-time traffic, besides the best-effort scheduling service they share with the DOCSIS 1.0 CMs on the same upstream channel.

The DOCSIS 1.1 upstream cable access channel carries variable-length MAC frames. In spite of the variable-length nature of the MAC frames, the DOCSIS 1.1 CMTS grant scheduler is theoretically capable of providing a zero jitter TDMA-like environment for voice grants on the Upstream. Whenever the grant scheduler detects that the deadline of any future voice grant will be violated by the insertion of a non-voice grant, it fragments the non-voice grant up to the future voice grant boundary. Thus the voice grants see a zero shift from the assigned periodic grant position.

However, such grant fragmentation might not always be possible when the CMTS supports DOCSIS 1.0 CMs along with DOCSIS 1.1 CMs on the same Upstream channel since DOCSIS 1.0 CM do not support fragmentation. For a mixed CM version upstream channel, the worst case voice grant jitter seen by the DOCSIS 1.1 CMs, is when a DOCSIS 1.0 CM is given a grant for an unfragmented maximum sized MAC frame just before the designated voice grant slot of the DOCSIS 1.1 CM.

The maximum Voice grant jitter experienced by the DOCSIS 1.1 CMs is a function of the physical layer characteristics of the Upstream Channel. For 10.24Mbps and 5.12Mbps upstream channels, the impact of having fragmenting and non-fragmenting CM's on the same channel is almost undetectable. On smaller channels, the benefit of fragmentation is far greater and the jitter induced by non-fragmenting DOCSIS 1.0 CMs is greater.

Thus, properly engineered networks can support voice even when mixing DOCSIS 1.0 and DOCSIS 1.1 CMs.

This page intentionally left blank.

Appendix H Multiple Upstream Channels

This section is informational. In case of conflict between this section and any normative section of this specification, the normative section takes precedence.

Section 5.2 describes support for multiple upstream and multiple downstream channels within a DOCSIS domain. The permutations that a CM may see on the cable segment it is attached to include:

- single downstream and single upstream per cable segment
- single downstream and multiple upstreams per cable segment
- multiple downstreams and single upstream per cable segment
- multiple downstreams and multiple upstreams per cable segment

A typical application that will require one upstream and one downstream per CM is web browsing. Web browsing tends to have asymmetrical bandwidth requirements that match closely to the asymmetrical bandwidth of DOCSIS.

A typical application that will require access to one of multiple upstreams per CM is IP Telephony. IP Telephony tends to have symmetrical bandwidth requirements. If there is a large concentration of CMs in a geographical area all served by the same fiber node, more than one upstream may be required in order to provide sufficient bandwidth and prevent call blocking.

A typical application that will require access to one of multiple downstreams per CM is IP streaming video. IP streaming video tends to have extremely large downstream bandwidth requirements. If there is a large concentration of CMs in a geographical area all served by the same fiber node, more than one downstream may be required in order to provide sufficient bandwidth and to deliver multiple IP Video Streams to multiple CMs.

A typical application that will require multiple downstreams and multiple upstreams is when the above applications are combined, and it is more economical to have multiple channels than it is to physically subdivide the HFC network.

The role of the CM in these scenarios would be to be able to move between multiple upstreams and between multiple downstreams. The role of the CMTS would be to manage the traffic load to all attached CMs, and balance the traffic between the multiple upstreams and downstreams by dynamically moving the CMs based upon their resource needs and the resources available.

This appendix looks at the implementation considerations for these cases. Specifically, the first and last application are profiled. These examples are meant to illustrate one topology and one implementation of that topology.

H.1 Single Downstream and Single Upstream per Cable Segment

This section presents an example of a single downstream channel and four upstream channels. In Figure H-1, the four upstream channels are on separate fibers serving four geographical communities of modems. The CMTS has access to the one downstream and all four upstreams, while each CM has access to the one downstream and only one upstream.

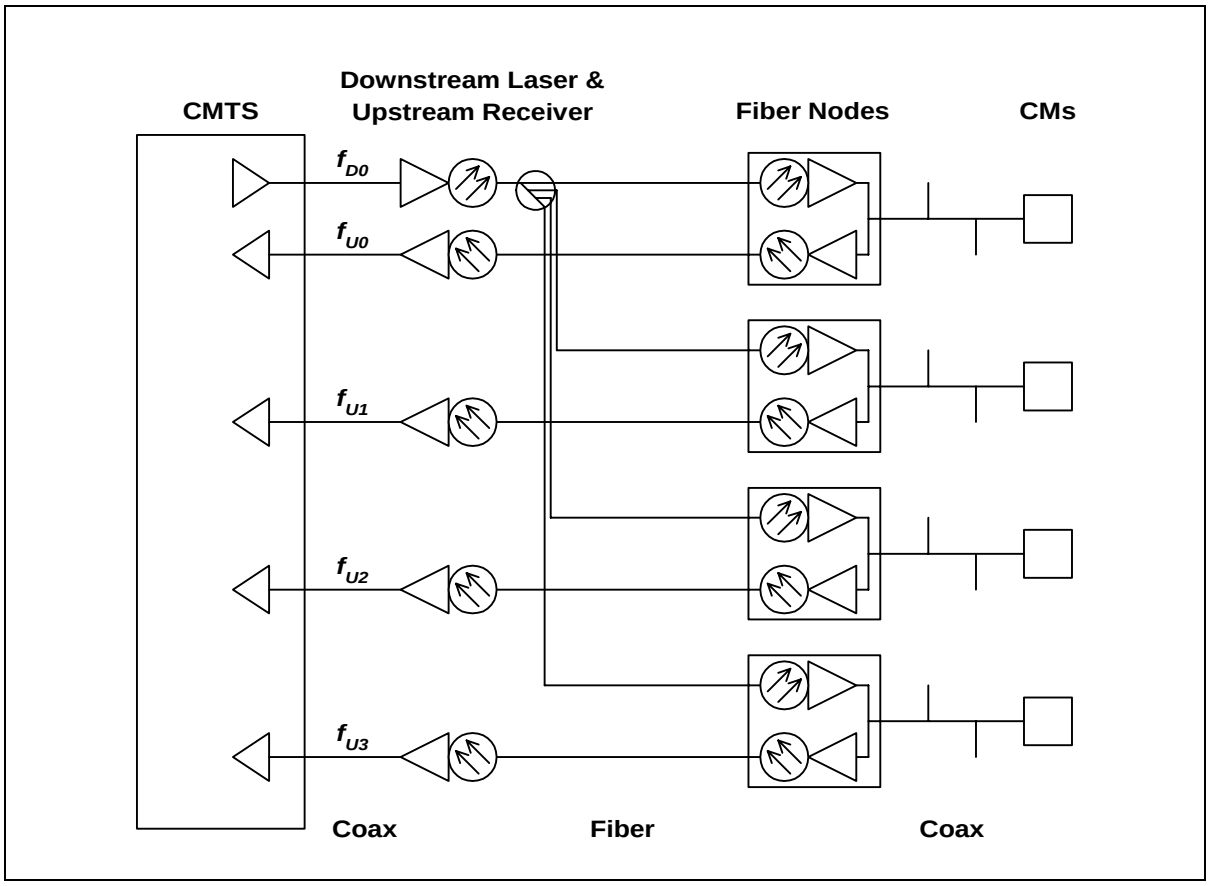


Figure H-1. Single Downstream and Single Upstream Channels per CM

In this topology, the CMTS transmits Upstream Channel Descriptors (UCDs) and MAPs for each of the four upstream channels related to the shared downstream channel.

Unfortunately, each CM cannot determine which fiber branch it is attached to because there is no way to convey the geographical information on the shared downstream channel. At initialization, the CM randomly picks a UCD and its corresponding MAP. The CM then chooses an Initial Maintenance opportunity on that channel and transmits a Ranging Request.

The CMTS will receive the Ranging Request and will redirect the CM to the appropriate upstream channel identifier by specifying the upstream channel ID in the Ranging Response. The CM MUST then use the channel ID of the Ranging Response, not the channel ID on which the Ranging 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. If the CM ever has to redo initial maintenance, it may start with its previous known UCD instead of choosing one at random.

A number of constraints are imposed by this topology:

- All Initial Maintenance opportunities across all fiber nodes must be aligned. When the CM chooses a UCD to use and then subsequently uses the MAP for that channel, the CMTS must be prepared to receive a Ranging Request at that Initial Maintenance opportunity. Note that only the initialization intervals must be aligned. Once the CM is successfully ranged on an upstream channel, its activities need only be aligned with other users on the same upstream channel. In Figure H-1, ordinary data transmission and requests for bandwidth may occur independently across the four upstream channels.
- All of the upstream channels on different nodes should operate at the same frequency or frequencies unless it is known that no other upstream service will be impacted due to a CM transmission of a Ranging Request on a "wrong" frequency during an Initial Maintenance opportunity. If the CM chooses an upstream channel descriptor arbitrarily, it could transmit on the wrong frequency if the selected UCD applied to an upstream channel on a different fiber node. This could cause initial maintenance to take longer. However, this might be an acceptable system trade-off in order to keep spectrum management independent between cable segments.
- All of the upstream channels may operate at different symbol rates. However, there is a trade-off involved between the time it takes to acquire ranging parameters and flexibility of upstream channel symbol rate. If upstream symbol rates are not the same, the CMTS would be unable to demodulate the Ranging Request if it was transmitted at the wrong symbol rate for the particular upstream receiver of the channel. The result would be that the CM would retry as specified in the RFI specification and then would eventually try other upstream channels associated with the currently used downstream. Increasing the probability of attempting ranging on multiple channels increases CM initialization time but using different symbol rates on different fiber nodes allows flexibility in setting the degree of burst noise mitigation.
- All Initial Maintenance opportunities on different channels may use different burst characteristics so that the CMTS can demodulate the Ranging Request. Again, this is a trade-off between time to acquire ranging and exercising flexibility in setting physical layer parameters among different upstream channels. If upstream burst parameters for Initial Maintenance are not the same, the CMTS would be unable to demodulate the Ranging Request if it was transmitted with the wrong burst parameters for the particular channel. The result would be that the CM would retry the Ranging Request as specified in the RFI specification and then would eventually try other upstream channels associated with the currently used downstream. Increasing the probability of attempting ranging on multiple channels increases CM initialization time but using different burst parameters for Initial Maintenance on different fiber nodes allows the ability to set parameters appropriate for plant conditions on a specific node.

H.2 Multiple Downstreams and Multiple Upstreams per Cable Segment

This section presents a more complex set of examples of CMs which are served by several downstream channels and several upstream channels and where those upstream and downstream channels are part of one MAC domain. The interaction of Initial Maintenance, normal operation, and Dynamic Channel Change are profiled, as well as the impact of the multiple downstreams using synchronized or unsynchronized timestamps.

Synchronized timestamps refer to both downstream paths transmitting a time stamp that is derived from a common clock frequency and have common time bases. The timestamps on each downstream do not have to be transmitted at the same time in order to be considered synchronized.

H.2.1 Topologies

Suppose two downstream channels are used in conjunction with four upstream channels as shown in Figure H-2. In all three topologies, there are two geographical communities of modems, both served by the same two downstream channels. The difference in the topologies is found in their upstream connectivity.

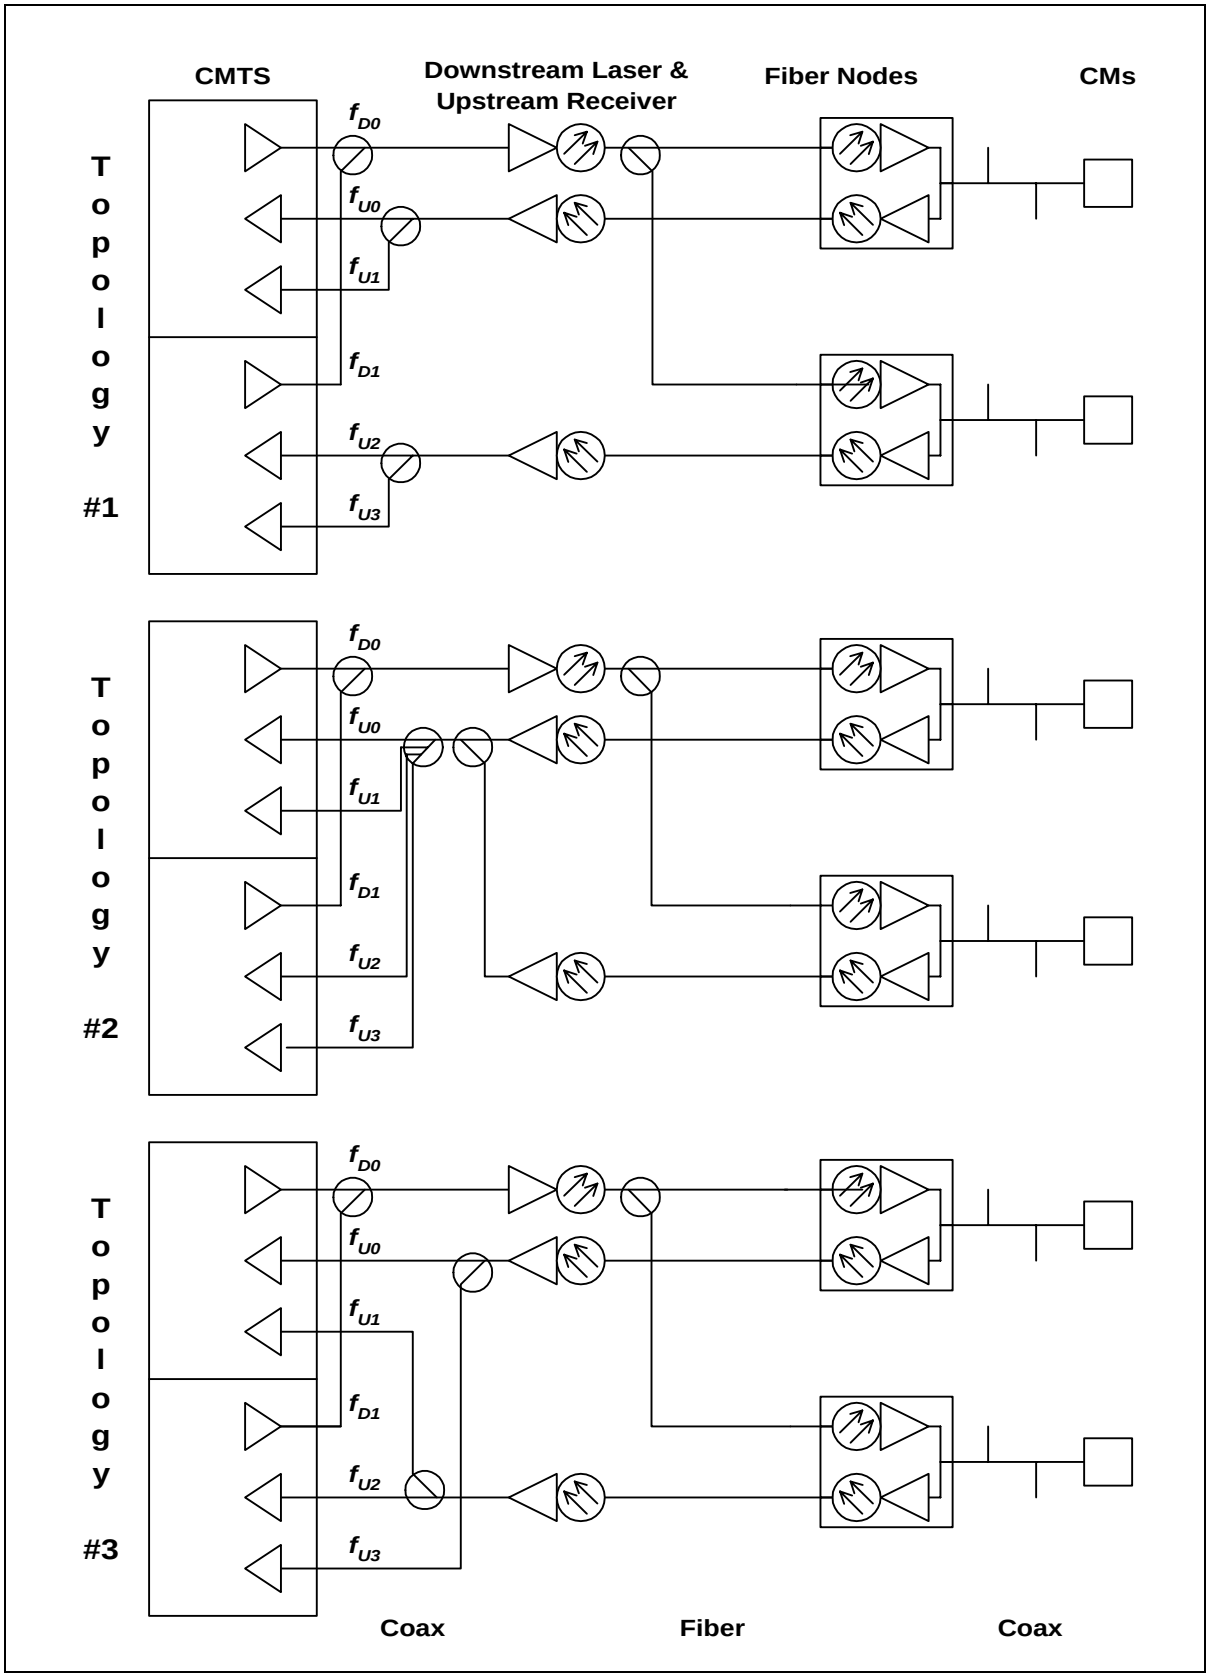


Figure H-2. Multiple Downstream and Multiple Upstream Channels per CM

Topology #1 has the return path from each fiber node connected to a dedicated set of upstream receivers. A CM will see both downstream channels, but only one upstream channel which is associated with one of the two downstream channels.

Topology #2 has the return path from each fiber node combined and then split across all upstream receivers. A CM will see both downstream channels and all four upstream channels in use with both downstream channels.

Topology #3 has the return path from each fiber node split and then sent to multiple upstream receivers, each associated with a different downstream channel. A CM will see both downstream channels, and one upstream channel associated with each of the two downstream channels.

Topology #1 is the typical topology in use. Movement between downstreams can only occur if the timestamps on both downstreams are synchronized. Topology #2 and Topology #3 are to compensate for downstreams which have unsynchronized timestamps, and allow movement between downstream channels as long as the upstream channels are changed at the same time.

The CMs are capable of single frequency receive and single frequency transmit.

H.2.2 Normal Operation

Table 8-1 lists MAC messages that contain Channel IDs.

Table 8-1. MAC Messages with Channel IDs

MAC Message	Downsteam Channel ID	Upstream Channel ID
UCD	Yes	Yes
MAP	No	Yes
RNG-REQ	Yes	No
RNG-RSP	No	Yes
DCC-REQ	Yes	Yes

With unsynchronized timestamps:

- Since upstream synchronization relies on downstream timestamps, each upstream channel must be associated with the time stamp of one of the downstream channels.
- The downstream channels should only transmit MAP messages and UCD messages that pertain to their associated upstream channels.

With synchronized timestamps:

- Since upstream synchronization can be obtained from either downstream channel, all upstreams can be associated with any downstream channel.
- All MAPs and UCDs for all upstream channels should be sent on all downstream channels. The UCD messages contains a Downstream Channel ID so that the CMTS can determine with the RNG-REQ message which downstream channel the CM is on. Thus the UCD messages on each downstream will contain different Downstream Channel IDs even though they might contain the same Upstream Channel ID.

H.2.3 Initial maintenance

When a CM performs initial maintenance, the topology is unknown and the timestamp consistency between downstreams is unknown. Therefore, the CM chooses either downstream channel and any one of the UCDs sent on that downstream channel.

In both cases:

- The upstream channel frequencies within a physical upstream or combined physical upstreams must be different.
- The constraints specified in Section H.1 apply.

H.2.4 Dynamic Channel Change

With unsynchronized timestamps:

- When a DCC-REQ is given, it must contain new upstream and new downstream frequency pairs that are both associated with the same timestamp.
- When the CM resynchronizes to the new downstream, it must allow for timestamp resynchronization without re-ranging unless instructed to do so with the DCC-REQ command.
- Topology #1 will support channel changes between local upstream channels present within a cable segment, but will not support changes between downstream channels. Topology #2 and #3 will support upstream and downstream channel changes on all channels within the fiber node as long as the new upstream and downstream channel pair are associated with the same timestamp.

With synchronized timestamps:

- Downstream channel changes and upstream channel changes are independent of each other.

Topology #1, #2, and #3 will support changes between all upstream and all downstream channels present within the cable segment.

Appendix I The Data-Over-Cable Spanning Tree Protocol

Section 3.1.2.1 requires the use of the spanning tree protocol on CMs that are intended for commercial use and on bridging CMTSs. This appendix describes how the 802.1d spanning tree protocol is adapted to work for data over cable systems.

I.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 [IEEE802.1Q/D] is employed in most bridged local area networks. This protocol was intended for private LAN use and requires some modification for cable data use.

I.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 I-1 illustrates the general topology.

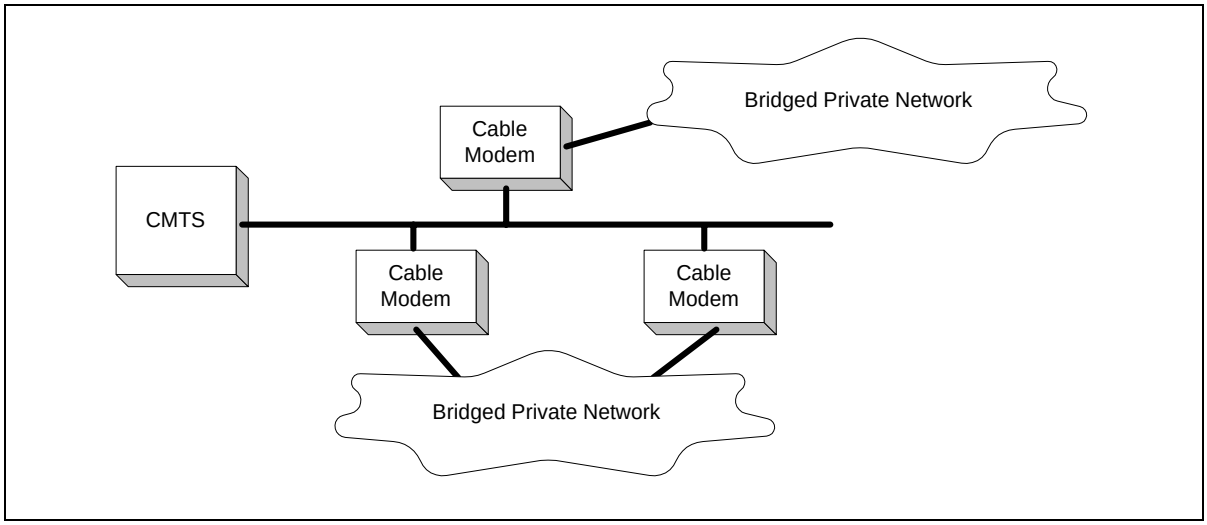


Figure I-1. Spanning Tree Topology

The task for the public spanning tree protocol, with reference to Figure I-1, is to:

- Isolate the private bridged networks from each other. If the two private networks merge spanning trees then 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 I-2:

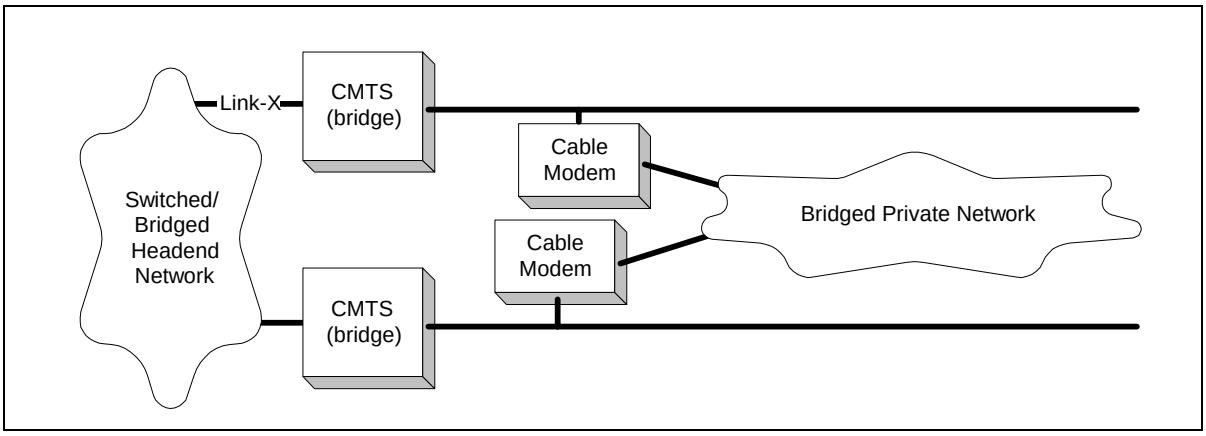


Figure I-2. Spanning Tree Across CMTSs

In Figure I-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, then it must be prevented by means external to spanning tree; for example, by using routers.

I.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 MUST be 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 MUST be 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 addresses.[1]

[1] 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.

- IEEE 802.1d BPDUs MUST be ignored and silently discarded.
- Topology Change Notification (TCN) PDUs MUST NOT be 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.
- CMTSs 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 bridged private subscriber network. Similarly, bridging CMTSs 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.

I.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:

Path_Cost = 1000 / Attached_LAN_speed_in_Mb/s

For CMs, this formula is adapted as:

Path_Cost = 1000 / (Upstream_symbol_rate * 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 I-1 provides the derived values.

Table I-1. CM Path Cost

Symbol Rate	Default Path Cost	
ksym/sec	QPSK	16QAM
160	3125	1563
320	1563	781
640	781	391
1280	391	195
2560	195	98

For CMTSs, this formula is:

Path_Cost = 1000 / (Downstream_symbol_rate * 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 J Error Codes and Messages

Please refer to SP-OSSIv1.1 [DOCSIS5], Appendix H for a complete list of error codes and messages.

This page intentionally left blank.

Appendix K DOCSIS Transmission and Contention Resolution

K.1 Introduction

This Appendix attempts to clarify how the DOCSIS transmission and contention resolution algorithms work. It has a few minor simplifications and a few assumptions, but should definitely help clarify this area of the specification.

This example has a few simplifications:

- It doesn't explicitly talk about packet arrivals while deferring or waiting for pending grants and is vague about sizing piggyback requests.
- The CM sends a Piggyback Request for the next frame in the last fragment and not inside one of the headers of the original frame.
- Much of this applies with concatenation, but it does not attempt to address all the subtleties of that situation.

It also has a few assumptions:

- It assumes that a Request always fits in any Request/Data region.
- When a piggyback request is sent with a contention data packet, the state machine only checks for the Grant to the Request and assumes the Data Ack for the contention data packet was supplied by CMTS.
- It probably assumes a few other things, but should be sufficient to get the basic point across.

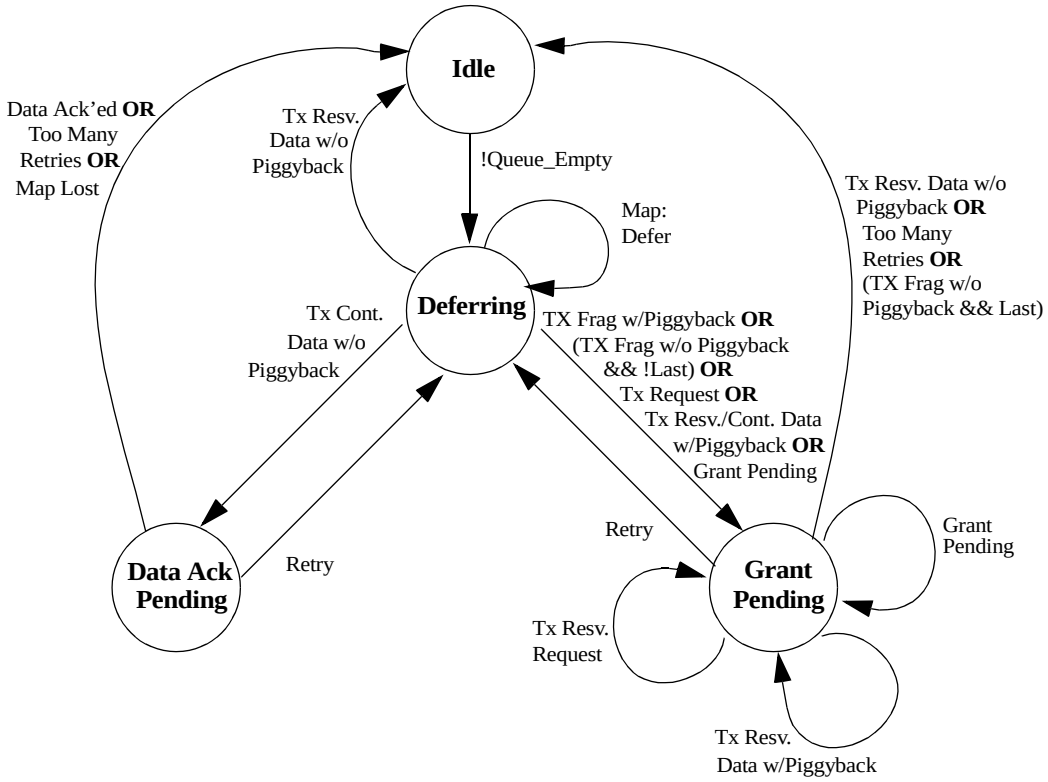


Figure K-1. Transmission & Deference State Transition Diagram

Variable Definitions:

Start	= Data Backoff Start field from Map “currently in effect”
End	= Data Backoff End field from Map “currently in effect”
Window	= Current backoff window
Random[n]	= Random number generator that selects a number between 0 and n-1
Defer	= Number of Transmit Opportunities to defer before transmitting
Retries	= Number of transmissions attempted without resolution
Tx_time	= Saved time of when Request or Request/Data was transmitted
Ack_time	= Ack Time field from current Map
Piggyback	= Flag set whenever a piggyback REQ is added to a transmit pkt
Queue_Empty	= Flag set whenever the data queue for this SID is empty
Lost_Map	= Flag set whenever a MAP is lost & we’re in state Data Ack Pending
my_SID	= Service ID of the queue that has a packet to transmit
pkt size	= Data packet size including MAC and physical layer overhead (including piggyback if used)
frag_size	= Size of the fragment
Tx_Mode	= {Full_Pkt; First_Frag; Middle_Frag; Last_Frag}
min_frag	= Size of the minimum fragment

State: Idle — Waiting for a Packet to Transmit

```
Window = 0;
Retries = 0;

Wait for !Queue_Empty;                  /* Packet available to transmit */

CalcDefer();
go to Deferring
```

State: Data Ack Pending — Waiting for Data Ack *only*

```
Wait for next Map;

if (Data Acknowledge SID == my_SID)     /* Success! CMTS received data packet */
    go to state Idle;
else if (Ack_time > Tx_time)            /* COLLISION!!! or Pkt Lost or Map Lost */
    {
    if (Lost_Map)
        go to state Idle;               /* Assume pkt was ack’ed to avoid sending duplicates */
    else
        Retry();
    }

stay in state Data Ack Pending;
```

State: Grant Pending — Waiting for a Grant

```
Wait for next Map;

while (Grant SID == my_SID)
    UtilizeGrant();

if (Ack_time > Tx_time)                 /* COLLISION!!!!! or Request denied/lost or Map Lost */
    Retry();
stay in state Grant Pending
```

State: Deferring — Determine Proper Transmission Timing & Transmit

```
if (Grant SID == my_SID)                          /* Unsolicited Grant */
     {
     UtilizeGrant();
     }
else if (unicast Request SID == my_SID)           /* Unsolicited Unicast Request */
     {
     transmit Request in reservation;
     Tx_time = time;

     go to state Grant Pending;
     }
else
     {
     for (each Request or Request/Data Transmit Opportunity)
          {
          if (Defer != 0)
               Defer = Defer - 1;                                    /* Keep deferring until Defer = 0 */
          else
               {
               if (Request/Data tx_op) and (Request/Data size >= pkt size)   /* Send data in contention */
                    {
                    transmit data pkt in contention;
                    Tx_time = time;

                    if (Piggyback)
                         go to state Grant Pending;
                    else
                         go to state Data Ack Pending;
                    }
               else                                                  /* Send Request in contention */
                    {
                    transmit Request in contention;
                    Tx_time = time;

                    go to state Grant Pending;
                    }
               }
          }
     }

Wait for next Map;
stay in state Deferring
```

Function: CalcDefer() — Determine Defer Amount

```
if (Window < Start)
     Window = Start;

if (Window > End)
     Window = End;

Defer = Random[2^Window];
```

Function: UtilizeGrant() — Determine Best Use of a Grant

```
if (Grant size >= pkt size)                    /* CM can send full pkt */
     {
     transmit packet in reservation;
     Tx_time = time;
     Tx_mode = Full_pkt

     if (Piggyback)
          go to state Grant Pending
     else
          go to state Idle;
     }
else if (Grant size < min_frag && Grant Size > Request size)/* Can't send fragment, but can send a Request */
     {
     transmit Request in reservation;
     Tx_time = time;

     go to state Grant Pending;
     }
else if (Grant size == 0)                      /* Grant Pending */
     go to state Grant Pending;
else
     {
     while (pkt_size > 0 && Grant SID == my_SID)
          {
          if (Tx_mode == Full_Pkt)
               Tx_mode = First_frag;
          else
               Tx_mode = Middle_frag;

          pkt_size = pkt_size - frag_size;
          if (pkt_size == 0)
               Tx_mode = Last_frag;

          if (another Grant SID == my_SID)       /* multiple grant mode */
               piggyback_size = 0
          else
               piggyback_size = pkt_size         /* piggyback mode */

          if (piggyback_size > 0)
               transmit fragment with piggyback request for remainder of packet in reservation
          else
               transmit fragment in reservation;
          }

     go to state Grant Pending;
```

Function: Retry()

```
Retries = Retries + 1;
if (Retries > 16)
    {
    discard pkt, indicate exception condition
    go to state Idle;
    }

Window = Window + 1;

CalcDefer();

go to state Deferring;
```

This page intentionally left blank.

Appendix L IGMP Example

Section 3.1.1 defines the requirements for CMTS and CM support of IGMP signaling. This Appendix provides an example CM passive-mode state machine for maintaining membership of a single multicast group.

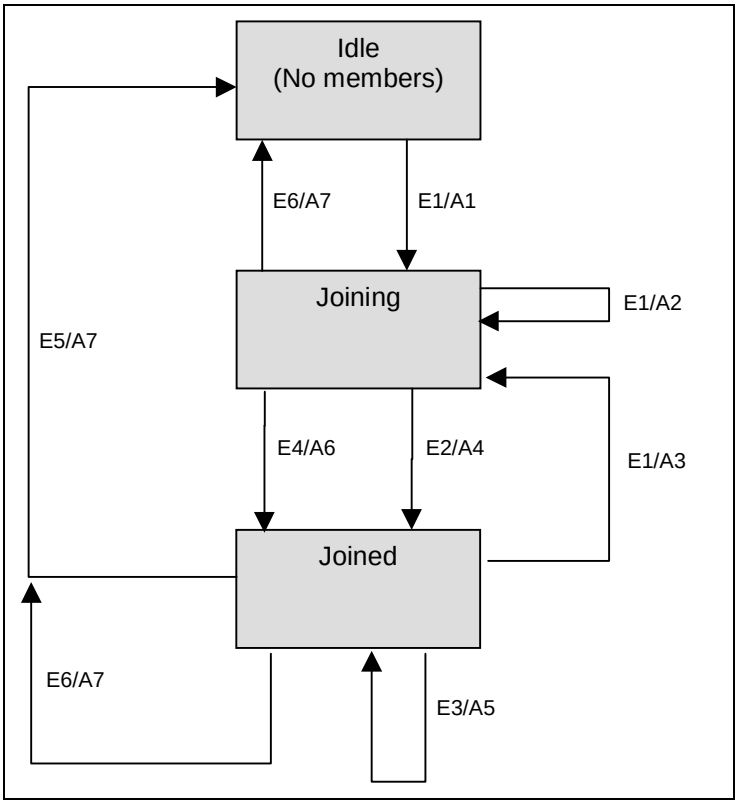


Figure L-1. IGMP Support – CM passive mode

Events

E1: MR received on CPE I/f
E2: M1 timer expired
E3: MQ received on RF I/f
E4: MR received on RF I/f
E5: M2 timer expired
E6: Auth Failure[1]

1. SA-MAP response returns an error code of 7 - "not authorized for requested downstream traffic flow"

Actions

A1: MQI= 125 sec; QRI = 10 sec; Start M1 timer with random value between 0 and 3 sec; start M2 timer = 2*MQI+QRI; start TEK machine, if necessary[2]; add multicast addr to multicast filter

A2: discard MR packet

A3: reset M2 timer = 2*MQI+QRI; start M1 timer with random value between 0 and 3 sec

A4: transmit MR on RF I/f; set I = current time

A5: recompute MQI = MAX(125, current time – I); set I = current time, forward MQ on CPE i/f

A6: cancel M1 timer

A7: delete multicast addr from multicast filter

[2.] If the multicast traffic is encrypted, then a TEK machine needs to be started to decrypt the encrypted multicast packets. To determine whether the multicast is encrypted, the CM makes a SA-MAP request to the CMTS to get the associated SAID of the multicast group address. If the SA-MAP response returns an SAID, then a TEK machine is started. No TEK machine is necessary, if the SA-MAP response indicates that the multicast traffic is not encrypted. The SA-MAP response may also indicate that the CM is not authorized to receive this multicast traffic. In which case, the CM terminates the multicast state machine and stops forwarding the multicast traffic.

Appendix M Unsolicited Grant Services

This appendix discusses the intended use of the Unsolicited Grant Service (UGS) and Unsolicited Grant Service with Activity Detection (UGS-AD) and includes specific examples.

M.1 Unsolicited Grant Service (UGS)

M.1.1 Introduction

Unsolicited Grant Service is an Upstream Flow Scheduling Service Type that is used for mapping constant bit rate (CBR) traffic onto Service Flows. Since the upstream is scheduled bandwidth, a CBR service can be established by the CMTS scheduling a steady stream of grants. These are referred to as unsolicited because the bandwidth is predetermined, and there are no ongoing requests being made.

The classic example of a CBR application of interest is Voice over Internet Protocol (VoIP) packets. Other applications are likely to exist as well.

Upstream Flow Scheduling Services are associated with Service Flows, each of which is associated with a single Service ID (SID). Each Service Flow may have multiple Classifiers. Each Classifier may be associated with a unique CBR media stream. Classifiers may be added and removed from a Service Flow. Thus, the semantics of UGS must accommodate single or multiple CBR media streams per SID.

For the discussion within this Appendix, a Subflow will be defined as the output of a Classifier. Since a VoIP session is identified with a Classifier, a Subflow in this context refers to a VoIP session.

M.1.2 Configuration Parameters

- Nominal Grant Interval
- Unsolicited Grant Size
- Tolerated Grant Jitter
- Grants per Interval

Explanation of these parameters and their default values are provided in Appendix C.

M.1.3 Operation

When a Service Flow is provisioned for UGS, the Nominal Grant Interval is chosen to equal the packet interval of the CBR application. For example, VoIP applications with 10 ms packet sizes will require a Nominal Grant Interval of 10 ms. The size of the grant is chosen to satisfy the bandwidth requirements of the CBR application and relates directly to the length of the packet.

When multiple Subflows are assigned to a UGS service, multiple grants per interval are issued. There is no explicit mapping of Subflows to grants. The multiple grants per interval form a pool of grants in which any subflow can use any grant.

It is assumed in this operational example the default UGS case of no concatenation and no fragmentation.

M.1.4 Jitter

Figure M-1 shows the relationship between Grant Interval and Tolerated Grant Jitter, and shows an example of jitter on subflows.

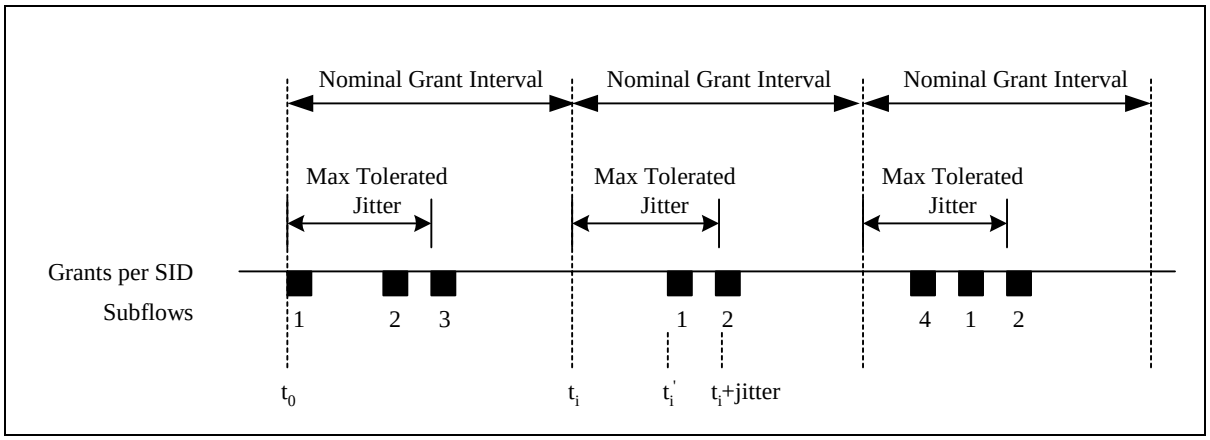


Figure M-1. Example Jitter with Multiple Grants per SID

For only one Grant per Interval, the Tolerated Grant Jitter is the maximum difference between the actual grant time (t_i') and the nominal grant time (t_i). For multiple Grants per Interval, the Tolerated Grant Jitter is the maximum difference between the actual time of the last grant in the group of grants and the nominal grant time (t_i). If the arrival of any grant is at t_i', then $t_i <= t_i' <= t_i$+jitter.

Figure M-1 demonstrates how a Subflow will be jittered even though the individual grants may not move from their relative position. During the first interval, three VoIP sessions are established, and they happen fall on the three grants. In the second interval, VoIP session 3 has been torn down. Since the CMTS does not know which Subflow is associated with which grant, it decides to remove the first grant. The remaining two calls shift to the other two grants. In the third interval, a new VoIP session 4 and a new grant have been added. The new call happens to fall on the new grant. The net effect is that the Subflows may move around within their jitter interval.

The advantage of a small jitter interval is that the VoIP receive jitter buffer may be kept small. The disadvantage is that this places a scheduling constraint on the CMTS.

The boundary of a Nominal Grant Interval is arbitrary and is not communicated between the CMTS and the CM.

Note: More dramatic events like the loss of a downstream MAP, or the frequency hopping of an upstream may cause subflows to jitter outside of this jitter window.

M.1.5 Synchronization Issues

There are two synchronization problems that occur when carrying CBR traffic such as VoIP sessions across a network. The first is a frequency mismatch between the source clock and the destination clock. This is managed by the VoIP application, and is beyond the scope of this specification. The second is the frequency mismatch between the CBR source/sinks, and the bearer channel that carries them.

Specifically, if the clock that generates the VoIP packets towards the upstream is not synchronized with the clock at the CMTS which is providing the UGS service, the VoIP packets may begin to accumulate in the CM. This could also occur if a MAP was lost, causing packets to accumulate.

When the CM detects this condition, it asserts the Queue Indicator in the Service Flow EH Element. The CMTS will respond by issuing an occasional extra grant so as to not exceed 1% of the provisioned bandwidth. (This

corresponds to a maximum of one extra grant every one hundred grants). The CMTS will continue to supply this extra bandwidth until the CM deasserts this bit.

A similar problem occurs in the downstream. The far end transmitting source may not be frequency synchronized to the clock which drives the CMTS. Thus the CMTS SHOULD police at a rate slightly higher than the exact provisioned rate to allow for this mismatch and to prevent delay buildup or packet drops at the CMTS.

M.2 Unsolicited Grant Service with Activity Detection (UGS-AD)

M.2.1 Introduction

Unsolicited Grant Service with Activity Detection (UGS-AD) is an Upstream Flow Scheduling Service Type. This section describes one application of UGS-AD which is the support for Voice Activity Detection (VAD). VAD is also known as Silence Suppression and is a voice technique in which the transmitting CODEC sends voice samples only when there is significant voice energy present. The receiving CODEC will compensate for the silence intervals by inserting silence or comfort noise equal to the perceived background noise of the conversation.

The advantage of VAD is the reduction of network bandwidth required for a conversation. It is estimated that 60% of a voice conversation is silence. With that silence removed, that would allow a network to handle substantially more traffic.

Subflows in this context will be described as active and inactive. Both of these states of within the MAC Layer QOS state known as Active.

M.2.2 MAC Configuration Parameters

The configuration parameters include all of the normal UGS parameters, plus:

- Nominal Polling Interval
- Tolerated Poll Jitter

Explanation of these parameters and their default values are provided in Appendix C.

M.2.3 Operation

When there is no activity, the CMTS sends polled requests to the CM. When there is activity, the CMTS sends Unsolicited Grants to the CM. The CM indicates the number of grants per interval which it currently requires in the active grant field of the UGSH in each packet of each Unsolicited Grant. The CM may request up to the maximum active Grants per Interval. The CM constantly sends this state information so that no explicit acknowledgment is required from the CMTS.

It is left to the implementation of the CM to determine activity levels. Implementation options include:

- Having the MAC layer service provide an activity timer per Classifier. The MAC layer service would mark a Subflow inactive if packets stopped arriving for a certain time, and mark a Subflow active the moment a new packet arrived. The number of Grants requested would equal the number of active Subflows.
- Having a higher layer service entity such as an embedded media client which indicates activity to the MAC layer service.

When the CM is receiving polled requests and it detects activity, the CM requests enough bandwidth for one Grant per Interval. If activity is for more than one Subflow, the CM will indicate this in the active grant field of the UGSH beginning with the first packet it sends.

When the CM is receiving Unsolicited Grants, then detects new activity, and asks for one more grant, there will be a delay in time before it receives the new grant. During that delay, packets may build up at the CM. When the new Unsolicited Grant is added, the CMTS will burst extra Grants to clear out the packet buildup.

When the CM is receiving Unsolicited Grants, then detects inactivity on a Subflow and asks for one less grant, there will be a delay in time before the reduction in Grants occurs. If there has been any build up of packets in the upstream transmit queue, the extra grants will reduce or empty the queue. This is fine, and keeps system latency low. The relationship of which Subflow is getting which specific grant will also change. This effect appears as low frequency jitter that the far end must manage.

When the CM is receiving Unsolicited Grants and detects no activity on any of its Subflows, it will send one packet with the active grants field of the UGSH set to zero grants, and then cease transmission. The CMTS will switch from UGS mode to Real Time Polling mode. When activity is again detected, the CM sends a request in one of these polls to resume delivery of Unsolicited Grants. The CMTS ignores the size of the request and resumes allocating Grant Size grants to the CM.

It is not necessary for the CMTS to separately monitor packet activity since the CM does this already. Worst case, if the CMTS misses the last packet which indicated zero grants, the CMTS and CM would be back in sync at the beginning of the next talk spurt. Because of this scenario, when the CM goes from inactive to active, the CM must be able to restart transmission with either Polled Requests or Unsolicited Grants.

M.2.4 Example

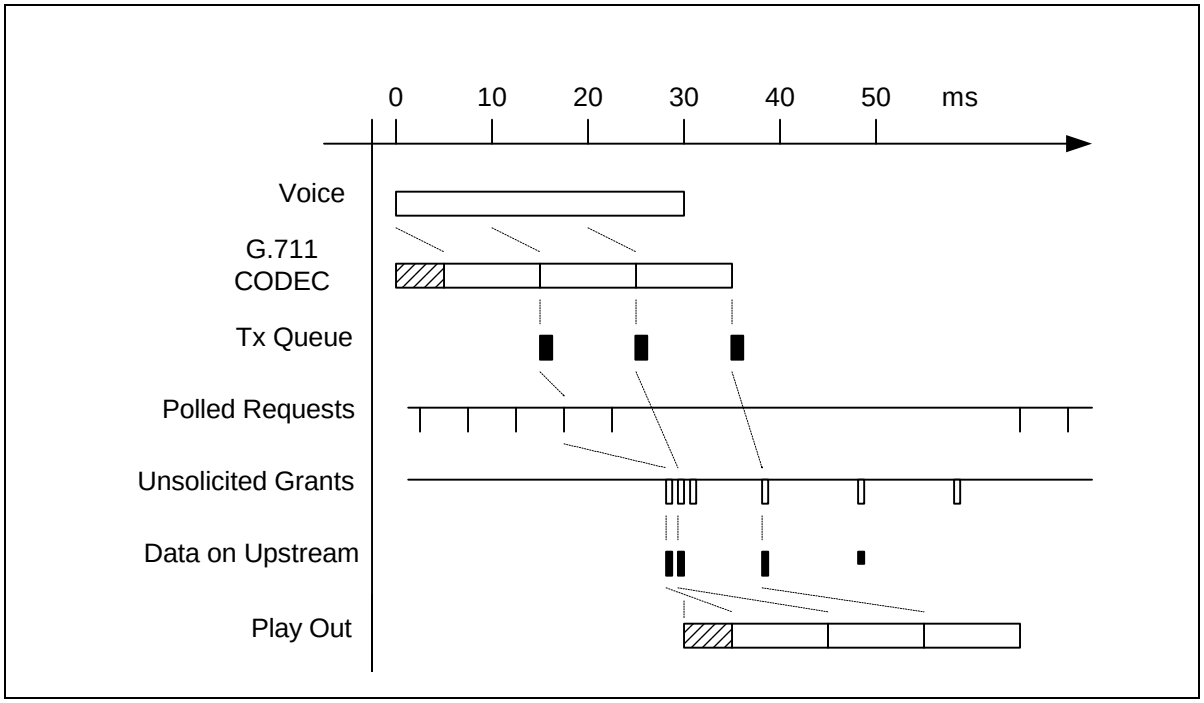


Figure M-2. VAD Start-Up and Stop

Figure M-2 shows an example of a single G.711 (64 kbps) voice call with a packet size of 10 ms, and a receive jitter buffer that requires a minimum of 20 ms of voice (thus 2 packets) before it will begin playout.

Assume voice begins at time zero. After a nominal processing delay and a 10 ms packetization delay, the DSP CODEC generates voice packets which are then transferred to the upstream transmit queue. The next Polled Request is used which results in the start of the Unsolicited Grants some time later. Additional Unsolicited Grants are immediately issued to clear out the upstream queue.

These packets traverse the network and arrive at the receive jitter buffer. The 20 ms minimum jitter buffer is met when the second packet arrives. Because the packets arrived close together, only an additional few milliseconds of latency has been added. After a nominal processing delay, playout begins.

When the voice spurt ends, the CM sends one remaining packet with no payload and with the active grants field of the UGSH set to zero grants. Some time later, UGS stops, and Real Time Polling begins.

M.2.5 Talk Spurt Grant Burst

The extra burst of Unsolicited Grants when a flow becomes active is necessary because the jitter buffer at the receiving CODEC typically waits to have a minimum amount of voice samples before beginning the playout. Any delay between the arrival of these initial packets will add to the final latency of the phone call. Thus, the sooner the CMTS recognizes that the CM has packets to send and can empty the CM's buffer, the sooner those packet will reach the receiver, and the lower the latency that will be incurred in the phone call.

It is an indeterminate problem as to how many grants must be burst. When the CM makes its request for an additional grant, one voice packet has already accumulated. The CM has no idea how many extra grants to request as it has no idea of the round trip response time it will receive from the CMTS, and thus how many packets may accumulate. The CMTS has a better idea, although it does not know the far end jitter buffer requirements.

The solution is for the CMTS to choose the burst size, and burst these grants close together at the beginning of the talk spurt. This occurs when moving from Real Time Polling to UGS, and when increasing the number of UGS Grants per Interval.

A typical start-up latency that will be introduced by the Request to Grant response time is shown in Table M-1.

Table M-1. Example Request to Grant Response Time

Variable		Example Value	
1.	The time taken from when the voice packet was created to the time that voice packet arrives in the CM upstream queue.	0 - 1	ms
2.	The time until a polled request is received. The worst case time is the Polled Request Interval.	0 - 5	ms
3.	The Request-Grant response time of the CMTS. This value is affected by MAP length and the number of outstanding MAPS.	5 - 15	ms
4.	The round trip delay of the HFC plant including the downstream interleaving delay.	1 - 5	ms
	Total	**6 - 26**	**ms**

This number will vary between CMTS implementations, but a reasonable number of extra grants to expect from the example above would be:

Table M-2. Example Extra Grants for New Talk Spurts

UGS Interval	Extra Grants for New Talk Spurts
10 ms	2
20 ms	1
30 ms	0

Once again it is worth noting that the CMTS and CM cannot and do not associate individual Subflows with individual grants. That means that when current Subflows are active and a new Subflow becomes active, the new Subflow will immediately begin to use the existing pool of grants. This potentially reduces the start up latency of new talk spurts, but increases the latency of the other Subflows. When the burst of grants arrives, it is shared with all the Subflows, and restores or even reduces the original latency. This is a jitter component. The more Subflows that are active, the less impact that adding a new Subflow has.

M.2.6 Admission Considerations

Note that when configuring the CMTS admission control, the following factors must be taken into account.

VAD allows the upstream to be over provisioned. For example, an upstream that might normally handle 24 VoIP sessions might be over provisioned as high as 36 (50%) or even 48 (100%). Whenever there is over provisioning, there exists the statistical possibility that all upstream VoIP sessions may become active. At that time, the CMTS may be unable to schedule all the VoIP traffic. Additionally, the talk spurt grant bursts would be stretched out. CM implementations of VAD should recognize this possibility, and set a limit as to how many packets they will allow to accumulate on its queue.

Occasional saturation of the upstream during VAD can be eliminated by provisioning the maximum number of permitted VoIP sessions to be less than the maximum capacity of the upstream with all voice traffic (24 in the previous example). VAD would cause the channel usage to drop from 100% to around 40% for voice, allowing the remaining 60% to be used for data and maintenance traffic.

Appendix N European Specification Additions

This section applies to the second technology option referred to in Section 1.1 (Scope). For the first option, refer to Section 2, Section 4 and Section 5.

This appendix describes the physical layer specifications required for what is generally called EuroDOCSIS cable-modems. This is an optional appendix and in no way affects certification of North American, DOCSIS 1.1 modems.

The numbering of the paragraphs has been made so that the suffix after the N refers to the part of the specification which has changed. As a consequence some paragraphs are missing in this appendix, because no change is required.

N.1 Scope and Purpose

No change 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 analogue 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 analogue 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 Appendix 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 analogue 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 analogue 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 (analogue 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 analogue modulated carriers	Not greater than – 57 dBc within the design bandwidth (Note 6 a)
Composite second-order distortion for analogue 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 analogue video carrier level at the system outlet, inclusive of above signal level variation	77 dBµV (Note 6 c)
Lowest analogue video carrier level at the system outlet, inclusive of above signal level variation	60 dBµV (Note 6 d)

NOTE 1 – Transmission is from the headend combiner to the CM input at the customer location.

NOTE 2 – For measurements above, the normal downstream operating frequency band (except hum), impairments are referenced to the highest-frequency PAL/SECAM carrier level.

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.

NOTE 4 – This presumes that the digital carrier is operated at analogue peak carrier level. When the digital carrier is operated below the analogue peak carrier level, this C/N may be less.

NOTE 5 – Measurements methods are defined in [CENELEC 50083-7]

NOTE 6 – For SECAM systems the following values apply

a) Not greater than -52 dBc within the design bandwidth

b) Not greater than -52 dBc within the design bandwidth

c) 74 dBµV

d) 57 dBµV

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 analogue video carrier level and will normally not exceed the analogue 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.3 Communication Protocols

No change required.

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 µs 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 band 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 Appendix N.4.2.7, N.4.2.8, and N.4.3.7. Maximum timing error and guard band may vary with CMTSes from different vendors. The term guard time is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.

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 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 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 Appendix N.4.2.6, Appendix 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, I_1 is the

MSB of the symbol map, Q_1 is the LSB for QPSK, and Q_0 is the LSB for 16QAM. Q_1 and I_0 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.

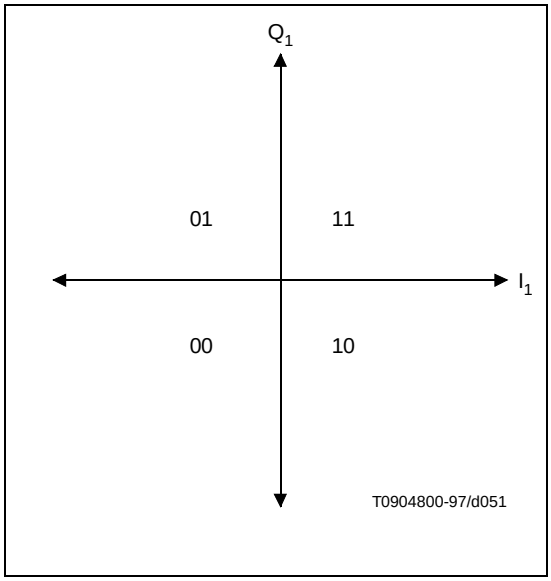


Figure N-1. QPSK Symbol Mapping

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

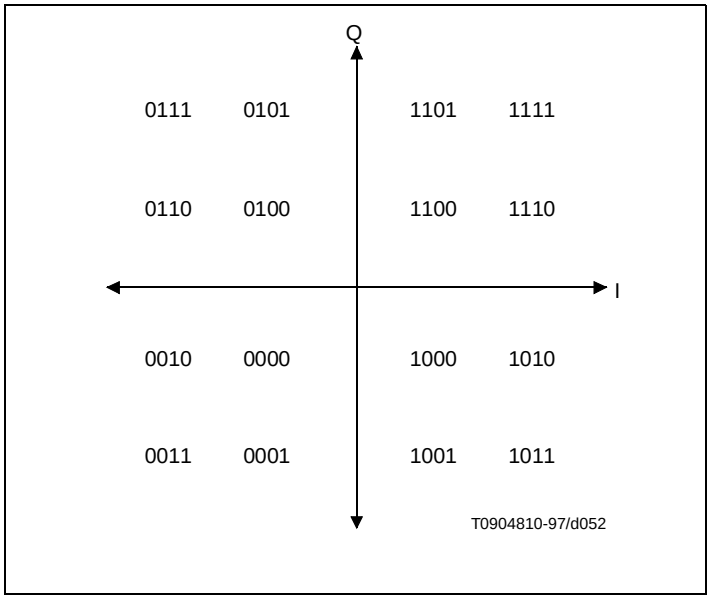


Figure N-2. 16QAM Gray-coded Symbol Mapping

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

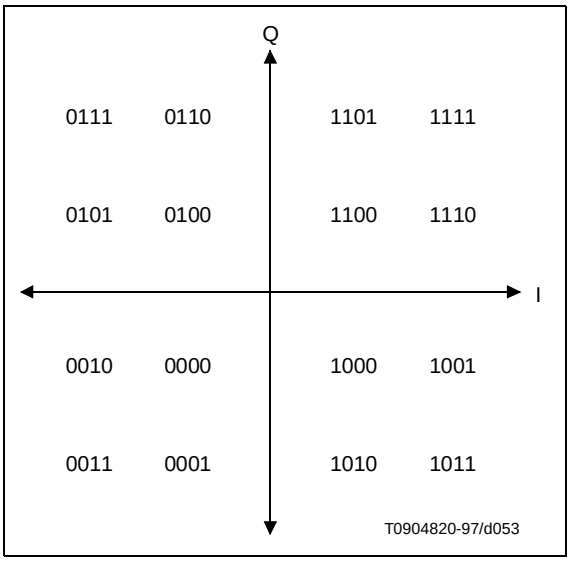


Figure N-3. 16QAM differential-coded symbol mapping

If differential quadrant encoding is enabled, then 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:

$g(x) = (x + \alpha^0)(x + \alpha^1)...(x + \alpha^{2T-1})$

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 Appendix N.4.2.11.1.2.

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 (Randomizer)

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 $x^{15} + x^{14} + 1$.

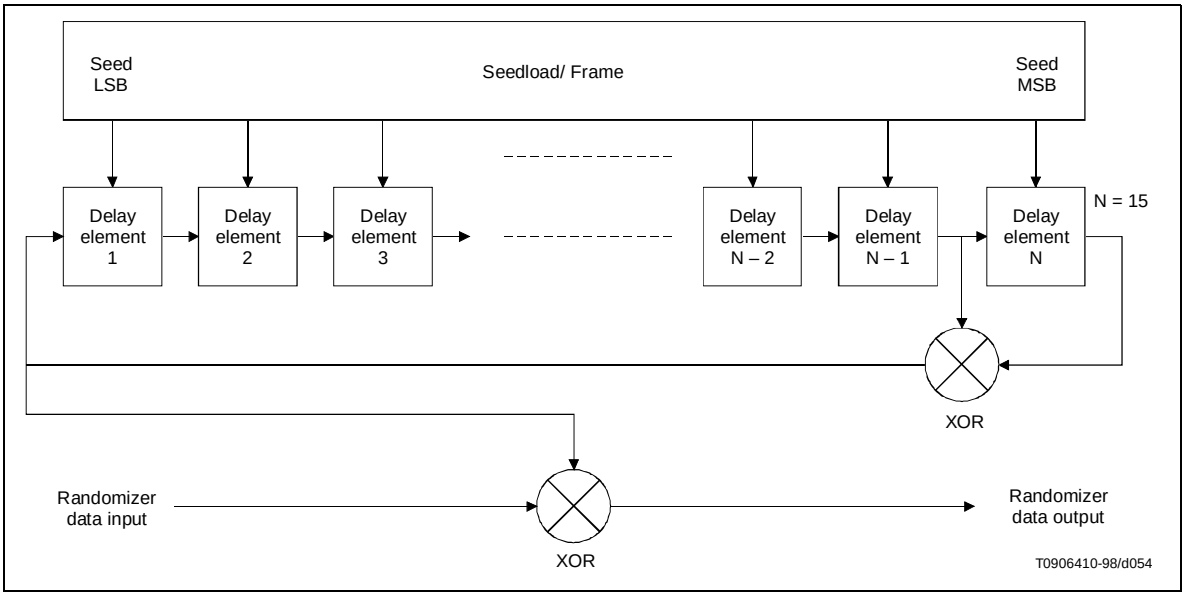


Figure N-4. Scrambler Structure

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 randomized and Reed-Solomon encoded.

The first bit of the Preamble Pattern is the first bit into the symbol mapper (see Figure N-9), and is I_1 in the first symbol of the burst (see Appendix 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 Transmit Pre-Equalizer

A transmit pre-equalizer of a linear equalizer structure, as shown in Figure N-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 to 4 samples per symbol, with a tap length longer than 8 symbols.

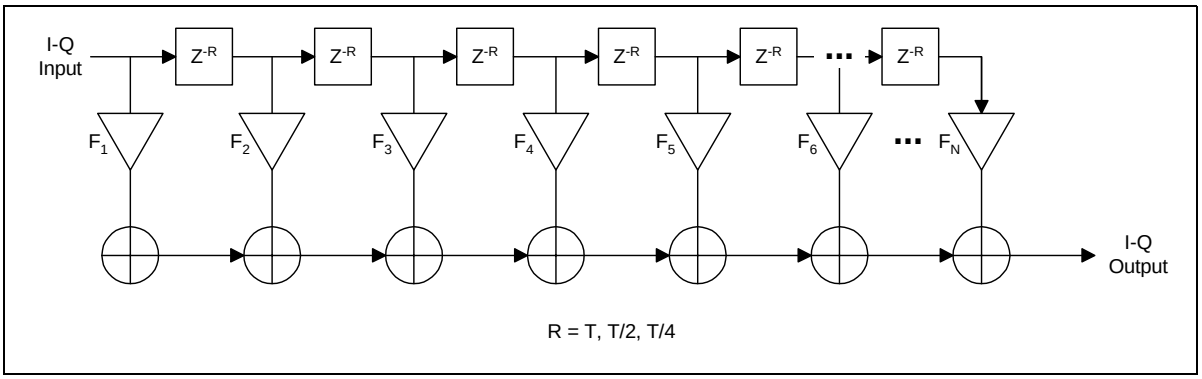


Figure N-5. Transmit Pre-Equalizer Structure

The RNG-RSP MAC message, (see Section 6.3.6.1) 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 6.3.8.1.1 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 though 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.

N.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 6.3.3, 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 Appendix N.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) (Note)	0 to 255
Guard Time	4 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 = ± 32 kHz; increment = 1 Hz; implement within ± 10 Hz
Ranging Offset	0 to (2^{16}-1), increments of 6.25 msec/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 µsec/64 = 0.09765625 µsec = 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 ± 0.25 µsec plus ± 1/2 symbol owing to resolution. The accuracy of CM burst timing of ± 0.25 µsec plus ± 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.

To provide backward interoperability with DOCSIS 1.0 and 1.1 equipment, when making a symbol rate change the CM MUST employ the following timing offsets when changing symbol rates. The offsets in the table correspond to the contribution of DOCSIS 1.0 and 1.1 legacy upstream receivers to changes in latency when

making symbol rate changes. The timing offset to apply is the difference between the entry in the table corresponding to the new symbol rate and the entry corresponding to the original symbol rate. The offsets are referenced to the center of the first symbol in the burst, which is the reference point for burst timing as stated in Section 4.2.8. Specification of these offsets is needed so that CMs apply uniform adjustments to their ranging offsets and so that CMTSs can appropriately handle CMs that apply these offsets when making symbol rate changes.

Symbol Rate	Timing Offset (in units of 1/64 time ticks referenced to 2.56 Msps)
2.56 Msps	0 (reference)
1.28 Msps	24
0.64 Msps	72
0.32 Msps	168
0.16 Msps	360

As an example, suppose a CM is on an upstream channel operating at 1.28 Msps. Now, suppose the UCD message from the CMTS changes the symbol rate of the channel to 0.32 Msps. The CM applies an additional timing offset of 168 – 24 = 144 to its ranging offset to compensate for this symbol rate change. The value of 144 is positive, and thus, the CM will add to its ranging offset so that it effectively transmits earlier by 144 units of 1/64 time ticks.

Furthermore, in changing symbol rates, if a CM has its own contribution to a change in latency, the CM MUST also compensate for this CM-specific latency difference. This is in addition to the offset applied from the values in the table above, which result from legacy CMTS upstream receiver contributions to changes in latency. The requirements for CM burst timing accuracy found earlier in this section, referenced to the symbol rate that is the lower of the original and the new symbol rate, apply after the symbol rate change with the required timing offsets above considered.

A CMTS that does not apply the same internal physical delay offsets as the legacy DOCSIS upstream CMTS receiver implementation is capable of receiving a CM burst after a symbol rate change in any of the following ways but is not limited necessarily to only these ways:

a) The CMTS may implement the internal physical delay offset, as specified in the above table.

b) The CMTS may implement an internal timing compensation based on the expected offset in the above table.

c) The CMTS may increase the guard time.

d) The CMTS may send an unsolicited RNG-RSP to each CM to adjust the delay offset. As discussed in Section 6.3.6, the CM is expected to be capable of adjusting its timing offset at any time with the accuracy specified within this section.

If Channel Frequency is to be changed, then 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 Appendix N.4.2.10.5 and Appendix N.4.2.10.6 within 100 msec from the beginning of the change.

If Output Power is to be changed by 1 dB or less, then 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, then 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:

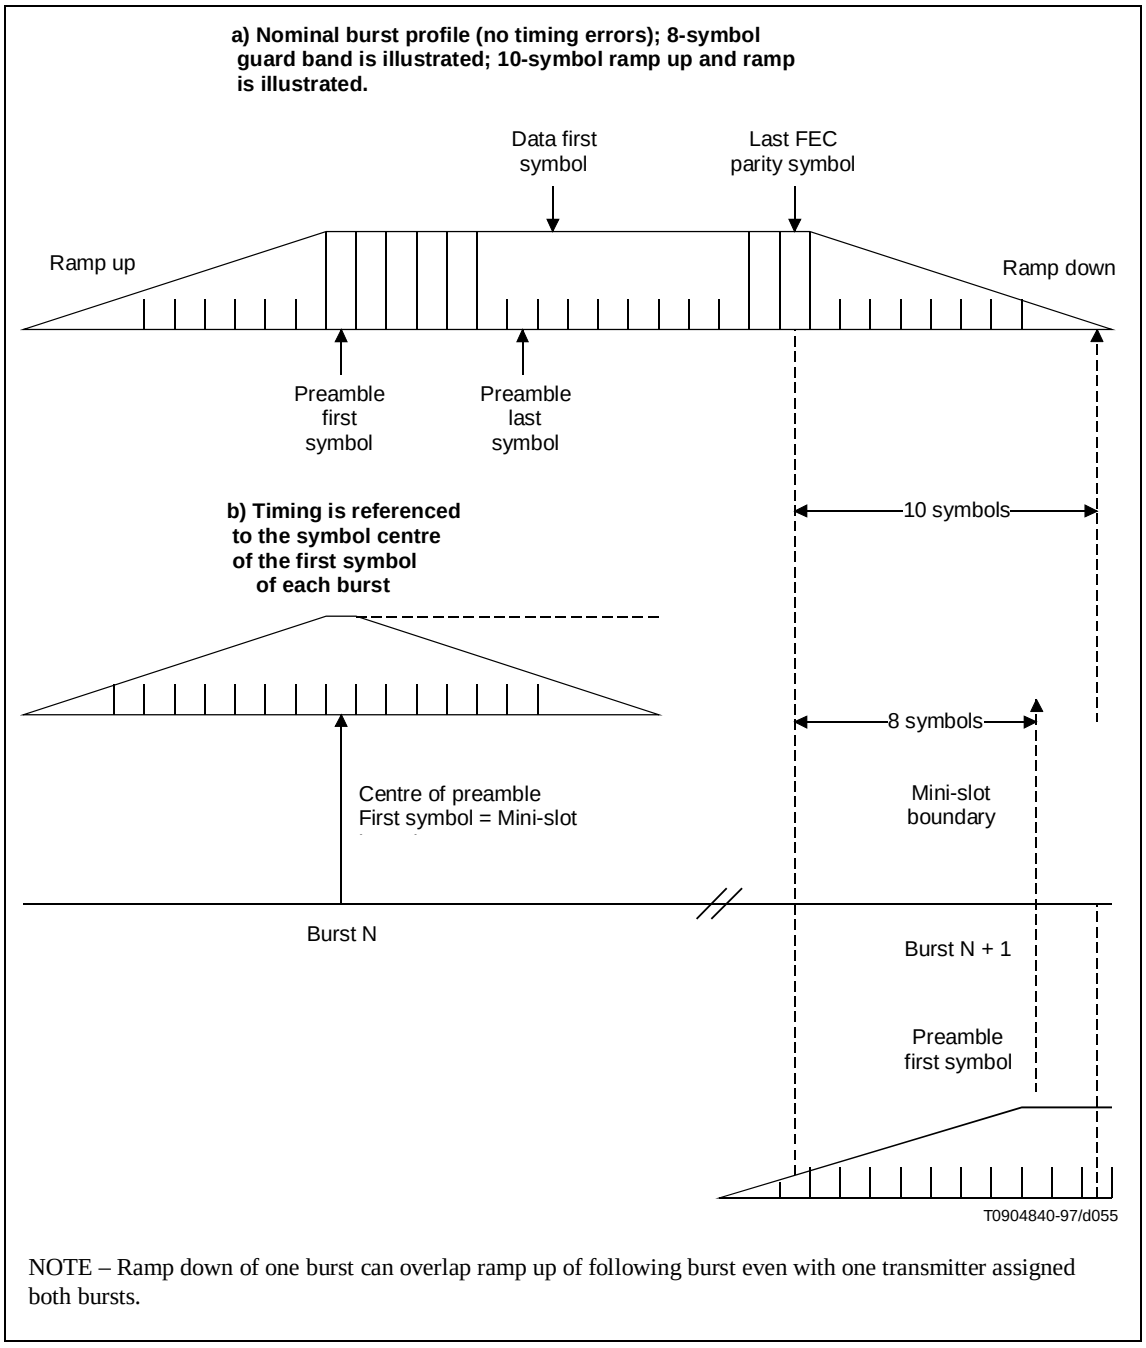
a) within 5 µsec from the beginning of a change of 1 dB or less; and

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.4.2.8 Burst timing convention

Figure N-6 illustrates the nominal burst timing.



NOTE – Ramp down of one burst can overlap ramp up of following burst even with one transmitter assigned both bursts.

Figure N-6. Nominal burst timing

Figure N-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 guard band shown.

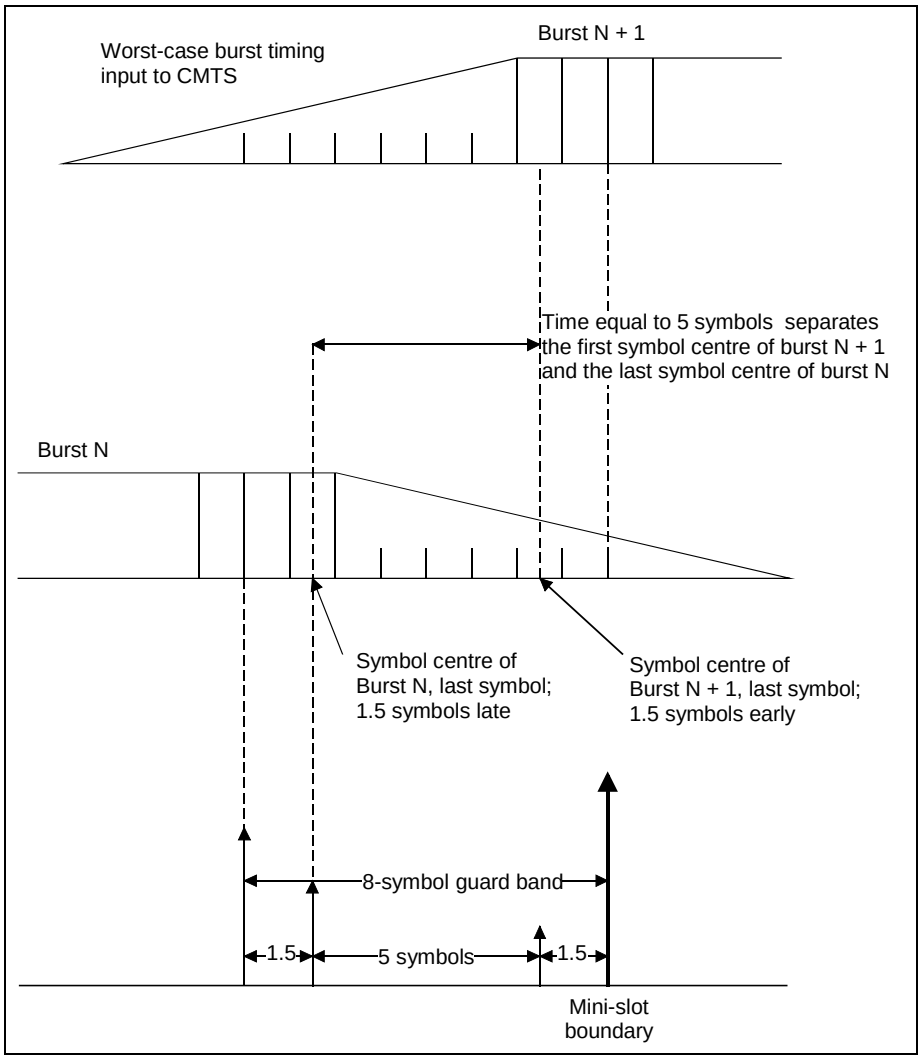


Figure N-7. Worst-case burst timing

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

N.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 9.2.4. Such adjustments MUST be within the ranges of tolerances described below.

N.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 (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.4.2.10 Fidelity requirements

N.4.2.10.1 Spurious emissions

The noise and spurious power MUST NOT exceed the levels given in Table N-8, Table N-9 and Table 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 µs at the 2.56 Msym/s rate, or as short as 200 µs 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 discretes) 87.5 to 108 MHz	 30 dBµV	 5 dBµV
CM Integrated Spurious Emissions Limits (all in 4.75 MHz, includes discretes) [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 specification 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 specification limits exclude three or fewer discrete spurs. Such spurs must not be greater than 20 dBµV.

N.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. 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.4.2.10.1.2 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.4.2.10.2 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 +115 dBµV 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 dBµV, down to a maximum change of 7 mV at 91 dBµV and below. This requirement does not apply to CM power-on and power-off transients.

The slew rate limitations of 2 µs need not be considered for DC transients of less than 7mV.

N.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_o), 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 wave form 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.

The CM MUST be capable of achieving a cluster SNR of at least 27 dB in the presence of the channel micro-reflections defined in Table 4-2. Since the table does not bound echo delay for the –30 dBc case, for testing purposes it is assumed that the time span of the echo at this magnitude is less than or equal to 1.5 µsec.

N.4.2.10.4 Filter distortion

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

N.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.

N.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.

N.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.

N.4.2.10.6 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.4.2.10.7 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.4.2.10.8 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.

N.4.2.11 Frame structure

Figure N-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 Appendix N.4.2.11.1.

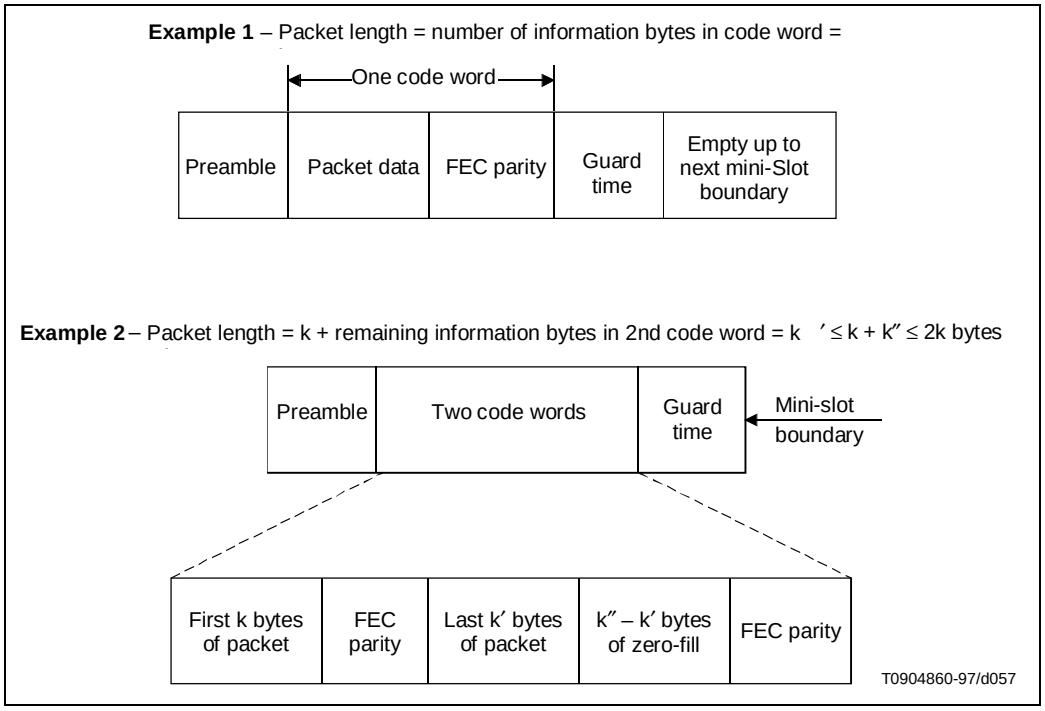


Figure N-8. Example frame structures with flexible burst length mode

N.4.2.11.1 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.4.2.11.1.1 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 min-slot in the grant, accounting for FEC parity and guard-time symbols.

N.4.2.11.1.2 Shortened last codeword

As shown in Figure N-8, 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 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″ – k′ zero-fill bytes are necessary to fill out a mini-slot allocation with $16 \le k'' \le k$ and $k' \le k''$. 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, then the CM shall not create this last shortened codeword.

N.4.2.12 Signal processing requirements

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

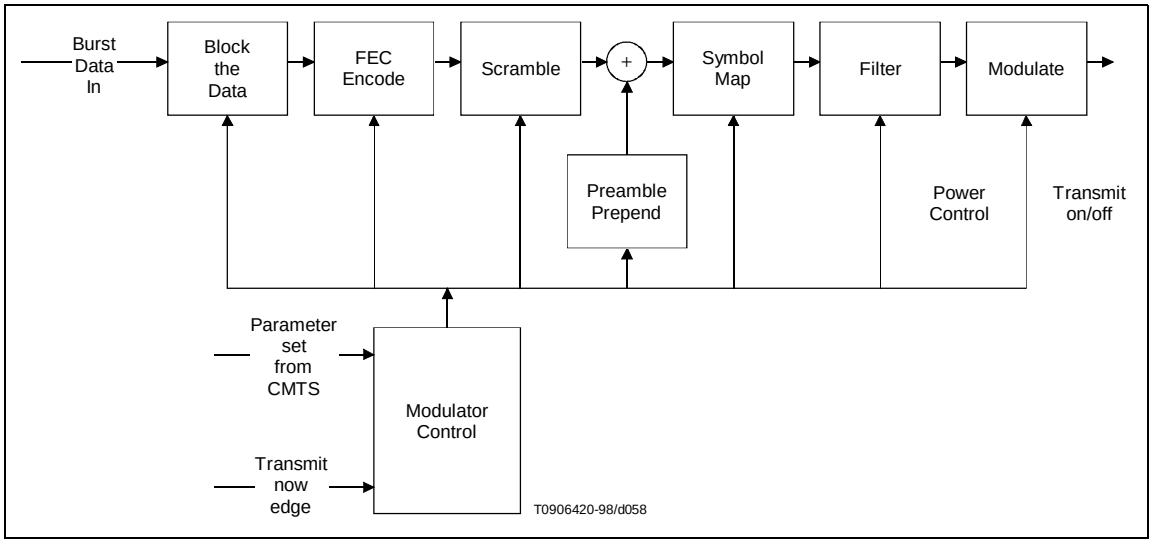


Figure N-9. Signal-processing sequence

	Packet Stream Input
	⇓
Block the Data	Separate Packet into Information Blocks (= data bytes in one code word)
	⇓
FEC Encode	FEC Encode (Reed-Solomon) each Information Block, using shortened codeword for last block if needed. FEC can be turned off.
	⇓
Scramble	Scramble (see Appendix N.4.2.4)
	⇓
Preamble Prepend	Preamble prepend Symbols
	⇓
Symbol Map	Symbol Map the Data Stream into Modulator Symbols
	⇓
Filter	Filter Symbol Stream for Spectral Shaping
	⇓
Modulate	Modulate at Precise Times (QPSK; 16QAM)
	⇓
	Output RF Wave-form Bursts

Figure N-10. TDMA upstream transmission processing

N.4.2.13 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 +77
640	±50 to +80
1280	+53 to +83
2560	+56 to +86

N.4.2.14 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)	+68 to +115 dBμV (16QAM) +68 to +118 dBμV (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	F connector per [ISO-169-24] (common with the input)

N.4.3 Downstream

N.4.3.1 Downstream protocol

The downstream PMD sublayer MUST conform to [EN 300 429].

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 µs/14 µs	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
Center Frequency (f_c)	112 to 858 MHz ± 30 kHz
Level	Adjustable over the range 110 to 121 dBµV
Modulation type	64QAM and 256QAM
Symbol rate (nominal) 64QAM 256QAM	 6.952 Msym/s 6.952 Msym/s
Nominal channel spacing	8 MHz
Frequency response 64QAM 256QAM	 ~15% square root raised cosine shaping ~15% square root raised cosine shaping
Total discrete spurious In-band (f_c ± 4 MHz)	< –57 dBc
In-band spurious and noise (f_c ± 4 MHz)	< –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 ± 4.0 MHz) to (f_c ±4.75 MHz)	< –58 dBc in 750 kHz.
Adjacent channel (f_c ± 4.75 MHz) to (f_c ±12 MHz)	< –60.6 dBc in 7.25 MHz, excluding up to 3 spurs, each of which must be < –60 dBc when each is measured with 10 kHz bandwidth.
Next adjacent channel (f_c ± 12 MHz) to (f_c ± 20 MHz)	Less than the greater of –63.7 dBc or 49.3 dBµV in 8 MHz, excluding up to three discrete spurs. The total power in the spurs must be < –60 dBc when each is measured with 10 kHz bandwidth.
Other channels (80 MHz to 1000 MHz)	< 49.3 dBµV in each 8 MHz channel, excluding up to three discrete spurs. The total power in the spurs must be < – 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	> 14 dB within an output channel up to 750 MHz; > 13 dB in an output channel above 750 MHz
Connector	F connector per[ISO-169-24]

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
Center Frequency	112 to 858 MHz ± 30 kHz
Level Range (one channel)	43 to 73 dBµV for 64QAM 47 to 77 dBµV 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)	< 90 dBµV
Input (load) Impedance	75 ohms
Input Return Loss	> 6 dB (85 to 862 MHz)
Connector	F connector per [ISO-169-24] (common with the output)

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 = 12, J = 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 (Es/No) of 25.5 dB or greater.

N.4.3.6.1.2 64QAM image rejection performance

Performance as described in Appendix N.4.3 MUST be met with analogue 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 Appendix N.4.3 MUST be met with digital signal at 0 dBc in the adjacent channels.

Performance as described in Appendix N.4.3 MUST be met with analogue signal at +10 dBc in the adjacent channels.

Performance as described in Appendix N.4.3, 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 (Es/No) as shown in Table N-16.

Table N-16. 256QAM CM BER performance

Input receive signal level	Es/No
47 dBµV to 54 dBµV	34.5 dB
> 54 to +77 dBµV	31.5 dB

N.4.3.6.2.2 256QAM image rejection performance

Performance as described in Appendix N.4.3 MUST be met with analogue 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 Appendix N.4.3 MUST be met with analogue or digital signal at 0 dBc in the adjacent channels.

Performance as described in Appendix N.4.3, with an additional 0.5-dB allowance, MUST be met with analogue signal at +10 dBc in the adjacent channels.

Performance as described in Appendix N.4.3, 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:

$|(N2 - N1)/ 10240000 - (T2 - T1)| < 500$ nsec

The jitter includes inaccuracy in timestamp value and the jitter in all clocks. The 500 ns 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 Appendix N.4.2.7 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 ≤ 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 Appendix 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-11 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-11. 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-12. The packet consists of a 4-byte MPEG Header, a pointer_field (not present in all packets) and the EuroDOCSIS Payload.

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

Figure N-12. Format of an MPEG Packet

N.5.3 MPEG Header for EuroDOCSIS 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 EuroDOCSIS 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	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 EuroDOCSIS Data-Over-Cable

The MPEG Payload portion of the MPEG Packet will carry the EuroDOCSIS 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 EuroDOCSIS Payload to fill any gaps between the EuroDOCSIS MAC frames. This value is chosen as an unused value for the first byte of the EuroDOCSIS 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 EuroDOCSIS 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 EuroDOCSIS 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-13 shows a MAC Frame that is positioned immediately after the pointer_field byte. In this case, pointer_field is zero, and the EuroDOCSIS decoder will begin searching for a valid FC byte at the byte immediately following the pointer_field.

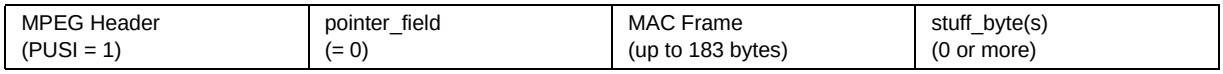


Figure N-13. Packet format where a MAC Frame immediately follows the pointer_field

Figure N-14 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.

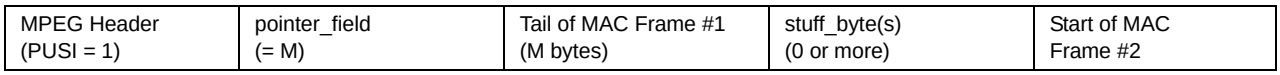


Figure N-14. Packet format with MAC Frame preceded by stuffing bytes

In order to facilitate multiplexing of the MPEG packet stream carrying EuroDOCSIS data with other MPEG-encoded data, the CMTS SHOULD NOT transmit MPEG packets with the EuroDOCSIS 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 EuroDOCSIS MAC sublayer which must also be preserved by any MPEG multiplexing operation.

Figure N-15 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.

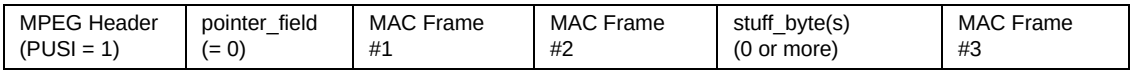


Figure N-15. Packet format showing multiple MAC frames in a single packet

Figure N-16 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-16. 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.

Appendix O 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 Service 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-I09-030730.

[DOCSIS5] Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification, SP-OSSIv1.1-I07-030730.

[DOCSIS6] Data-Over-Cable Service Interface Specifications, Cable Modem Telephony Return Interface Specification, SP-CMTRI-I01-970804.

[DOCSIS8] Data-Over-Cable Service Interface Specifications, Baseline Privacy Plus Interface Specification, SP-BPI+-I10-030730.

[EIA-S542] EIA Standard 542 (1997), “Cable Television Channel Identification Plan”, May 1997.

[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 EN 50081-1 Electromagnetic compatibility generic emission standard Part 1: Residential, commercial and light industry.

[EN 50082-1] CENELEC EN 50082-1 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”.

[EG 201 212] ETSI EG 201 212: “Electrical safety; Classification of interfaces for equipment to be connected to telecommunication networks”. This document is also available from CENELEC as ROBT-002.

[FCC15] Code of Federal Regulations, Title 47, Part 15, October 1998.

[FCC76] Code of Federal Regulations, Title 47, Part 76, October 1998.

[ID-IGMP] Fenner W., IGMP-based Multicast Forwarding ("IGMP Proxying"), IETF Internet Draft, http://www.ietf.org/internet-drafts/draft-ietf-magma-igmp-proxy-00.txt.

[IEEE802] IEEE Std 802-1990, Local and Metropolitan Area Networks: Overview and Architecture.

[IEEE802.1Q/D] IEEE Draft Standard 802.1Q/D4. Draft Standard for Virtual Bridged Local Area Networks. December 20, 1996.

[IMA] Internet Assigned Numbers Authority, Internet Multicast Addresses, http://www.isi.edu/in-notes/iana/assignments/multicast-addresses

[ISO-169-24] ISO-169-24 F connector, female, indoor

[ISO/IEC 8825-1] ISO/IEC 8825-1:2002 (ITU-T X.690) Information Technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER).

[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 (1995) | ISO/IEC 13818-1:1996, 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(03/93), 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 (3/93) - 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.

[PKTCBL-MGCP] PacketCable Specifications, Network-Based Call Signaling Protocol Specification, PKT-SP-EC-MGCP-I04-011221.

[PKT-DQOS] PacketCable Specifications, Dynamic Quality of Service Specification, PKT-SP-DQOS-I03-020116.

[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] Sollings, K., The TFTP Protocol (Revision 2), IETF RFC-1350, July, 1992.

[RFC-1493] Definitions of Managed Objects for Bridges. E. Decker, P. Langille, A. Rijsinghani, & 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-2210] Wroclawski, J., The Use of RSVP with the IETF Integrated Services, IETF RFC-2210, September, 1997.

[RFC-2211] Wroclawski, J., Specification of the Controlled-Load Network Element Service, IETF RFC-2211, September, 1997.

[RFC-2212] Shenker, S., Partridge, C., and Guerin, R., Specification of Guaranteed Quality of Service, IETF RFC-2212, September, 1997.

[RFC-2236] Fenner, W., Internet Group Management Protocol, Version 2, IETF RFC-2236, November 1997.

[RFC-2349] Malkin, G. and Harkin, A., TFTP Timeout Interval and Transfer Size Options, IETF RFC-2349, May 1998.

[RFC-2669] St. Johns, M., DOCSIS Cable Device MIB Cable Device Management Information Base for DOCSIS Compliant Cable Modems and Cable Modem Termination Systems, IETF RFC-2669, August 1999.

[RFC-2786] St. Johns, M., Diffie-Helman USM Key Management Information Base and Textual Convention, IETF RFC-2786, March, 2000.

[RFC-3046] Patrick, M., DHCP Relay Agent Information Option, IETF RFC-3046, January, 2001.

[RFC-3256] Jones, D. & Woundy, R., The DOCSIS Device Class DHCP Relay Agent Information Sub-option, IETF RFC-3256, April 2002.

[SCTE1] SCTE 22-1 2002, DOCSIS 1.0, Radio Frequency Interface Standard.

[SCTE2] SCTE 22-2 2002, DOCSIS 1.0, Baseline Privacy Interface Standard.

[SHA] NIST, FIPS PUB 180-1: Secure Hash Standard, April 1995.

[SMS] The Spectrum Management Application (SMA) and the Common Spectrum Management Interface (csmi), Time Warner Cable, December 24, 1995.

Appendix P Glossary

Active Service Flow
An admitted Service Flow from the CM to the CMTS which is available for packet transmission.

Address Resolution Protocol (ARP)
A protocol of the IETF for converting network addresses to 48-bit Ethernet addresses.

Admitted Service Flow
A Service Flow, either provisioned or dynamically signaled, which is authorized and for which resources have been reserved but is not active.

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.

Authorization Module
The authorization module is an abstract module that the CMTS can contact to authorize Service Flows and Classifiers. The authorization module tells the CMTS whether the requesting CM is authorized for the resources it is requesting.

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.

Bandwidth Allocation Map
The MAC Management Message that the CMTS uses to allocate transmission opportunities to CMs.

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.

Burst Error Second
Any Errored Second containing at least 100 errors.

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).

CCCM
CPE Controlled Cable Modem. Refer to the DOCSIS Cable Modem to Customer Premise Equipment Interface (CMCI) specification.

Classifier
A set of criteria used for packet matching according to TCP, UDP, IP, LLC, and/or 802.1P/Q packet fields. A classifier maps each packet to a Service Flow. A Downstream Classifier is used by the CMTS to assign packets to downstream service flows. An Upstream Classifier is used by the CM to assign packets to upstream service flows.

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.

Drop Cable
Coaxial cable that connects to a residence or service location from a directional coupler (tap) on the nearest coaxial feeder cable.

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.

ECN
See Engineering Change Notice.

ECO
See Engineering Change Order.

ECR
See Engineering Change Request.

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.

Engineering Change Notice
The final step in the procedure to change specifications.

Engineering Change Order
The second step in the procedure to change specifications. DOCSIS posts ECO to web site EC table and ECO page (with indication of ECO Comment Deadline). DOCSIS issues ECO announcement to DOCSIS-announce and working group mail lists (with indication of ECO Comment Deadline).

Engineering Change Request
The first step in the procedure to change specifications. DOCSIS issues ECR number, posts to web site EC table and ECR page. DOCSIS sends ECR to subject area working group mail list (and author).

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.

FDDI
See Fiber Distributed Data Interface.

Feeder Cable
Coaxial cables that run along streets within the served area and connect between the individual taps which serve the customer drops.

Fiber Distributed Data Interface (FDDI)
A fiber-based LAN standard.

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 Band
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 band should be at least the duration of five symbols plus the maximum system timing error.

Guard Time
The term guard time is similar to the guard band, except that it is measured from the end of the last symbol of one burst to the beginning of the first symbol of the preamble of an immediately following burst. Thus, the guard time is equal to the guard band – 1.

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.

High Return
A frequency division scheme that allows bi-directional traffic on a single coaxial cable. Reverse channel signals propagate to the headend above the downstream passband.

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.

IE
See Information Element.

IEEE
See Institute of Electrical and Electronic Engineers.

IETF
See Internet Engineering Task Force.

IGMP
See Internet Group Management Protocol.

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.

Internet Group Management Protocol (IGMP)
A network-layer protocol for managing multicast groups on the Internet

Impulse Noise
Noise characterized by non-overlapping transient disturbances.

Information Element
The fields that make up a MAP and define individual grants, deferred grants, etc.

Internet Protocol (IP)
An Internet network-layer protocol.

Interval Usage Code
A field in MAPs and UCDs to link burst profiles to grants.

IP
See Internet Protocol.

IUC
See Interval Usage Code.

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 6.1.2.2.

MAP
See Bandwidth Allocation Map.

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.

Mean Time to Repair (MTTR)
In cable television systems, the MTTR is the average elapsed time from the moment a loss of RF channel operation is detected up to the moment the RF channel operation is fully restored.

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.

Mid Split
A frequency division scheme that allows bi-directional traffic on a single coaxial cable. Reverse channel signals propagate to the headend from 5 to 108 MHz. Forward path signals go from the headend from 162 MHz to the upper frequency limit. The diplex crossover band is located from 108 to 162 MHz.

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 7.3.4.

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.

Multipoint Access
User access in which more than one terminal equipment is supported by a single network termination.

Multipoint Connection
A connection among more than two data network terminations.

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 framework of ISO standards for communication between different systems made by different vendors, in which the communications process is organized into seven different categories that are placed 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 through 4 deal with end-to-end communication between the message source and destination, and layers 3 through 1 deal with network functions.

Organizationally Unique Identifier (OUI)
A 3-octet IEEE assigned identifier that can be used to generate Universal LAN MAC addresses and Protocol Identifiers per ANSI/IEEE Std 802 for use in Local and Metropolitan Area Network applications.

OSI
See Open Systems Interconnection.

OUI
See Organization Unique Identifier.

Packet Identifier (PID)
A unique integer value used to identify elementary streams of a program in a single- or multi-program MPEG-2 stream.

Partial Grant
A grant that is smaller than the corresponding bandwidth request from the CM.

Payload Header Suppression
The suppression of the header in a payload packet. (*e.g.*, the suppression of the Ethernet header in forwarded packets)

Payload Unit Start Indicator (PUSI)

A flag in an MPEG header. A value of 1 indicates the presence of a pointer field as the first byte of the payload.

PHS
See Payload Header Suppression.

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.

Primary Service Flow
All CMs have a Primary Upstream Service Flow and a Primary Downstream Service Flow. They ensure that the CM is always manageable and they provide a default path for forwarded packets that are not classified to any other Service Flow

Program-Specific Information (PSI)
In MPEG-2, normative data necessary for the demultiplexing of Transport Streams and the successful regeneration of programs.

Program Stream
In MPEG-2, a multiplex of variable-length digital video and audio packets from one or more program sources having a common time-base.

Protocol
A set of rules and formats that determines the communication behavior of layer entities in the performance of the layer functions.

Provisioned Service Flow
A Service Flow that has been provisioned as part of the Registration process, but has not yet been activated or admitted. It may still require an authorization exchange with a policy module or external policy server prior to admission.

PSI
See Program-Specific Information.

QAM
See Quadrature Amplitude Modulation.

QoS Parameter Set
The set of Service Flow Encodings that describe the Quality of Service attributes of a Service Flow or a Service Class. (Refer to Appendix C.2.2.5)

QPSK
See Quadrature Phase-Shift Keying.

Quadrature Amplitude Modulation (QAM)
A method of modulating digital signals onto a radio-frequency carrier signal involving both amplitude and phase coding.

Quadrature Phase-Shift Keying (QPSK)
A method of modulating digital signals onto a radio-frequency 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; these documents can be accessed on the World Wide Web at 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.

Routing Information Protocol (RIP)
A protocol of the IETF for exchanging routing information about IP networks and subnets.

SAID
See Security Association Identifier.

Security Association Identifier
A Baseline Privacy security identifier between a CMTS and a CM.

Service Access Point (SAP)
The point at which services are provided by one layer, or sublayer to the layer immediately above it.

Service Class
A set of queuing and scheduling attributes that is named and that is configured at the CMTS. A Service Class is identified by a Service Class Name. A Service Class has an associated QoS Parameter Set.

Service Class Name
An ASCII string by which a Service Class may be referenced in modem configuration files and protocol exchanges.

Service Data Unit (SDU)
Information that is delivered as a unit between peer service access points

Service Flow
A MAC-layer transport service which:

- Provides unidirectional transport of packets from the upper layer service entity to the RF;
- Shapes, polices, and prioritizes traffic according to QoS traffic parameters defined for the Flow.

Service Flow Identifier (SFID)
An identifier assigned to a service flow by the CMTS. [32 bits]

Service Flow Reference
A message parameter in Configuration Files and Dynamic Service MAC messages used to associate Classifiers and other objects in the message with the Service Flow Encodings of a requested Service Flow.

Service Identifier (SID)
A Service Flow Identifier assigned by the CMTS (in addition to a Service Flow Identifier) to an Active or Admitted Upstream Service Flow. [14 bits]

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
6.25-microsecond time intervals that are the reference for upstream mini-slot definition and upstream transmission times.

Tilt
Maximum difference in transmission gain of a cable television system over a given bandwidth (typically the entire forward operating frequency range).

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.

Transmit On/Off Ratio
In multiple-access systems, the ratio between the signal powers sent to line when transmitting and when not transmitting.

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.

Trunk Cable
Cables that carry the signal from the headend to groups of subscribers. The cables can be either coaxial or fiber depending on the design of the system.

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.

Upstream Channel Descriptor (UCD)
The MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems.

This page intentionally left blank.

Appendix Q Acknowledgments

Rich Woundy of American Internet Corporation (now part of Cisco) led the Quality of Service working group that contributed additions to Sections 6 and 8 as well as Appendix C. Mike Patrick of Motorola and Burcak Beser of 3com were instrumental members of this group and contributed text and ideas for many portions of this specification.

Tom Quigley and Lisa Denney of Broadcom Corporation as well as John Ulm of Nortel Networks, Inc. Data Over Cable Division spearheaded the effort to incorporate fragmentation into the specification.

John Ulm, David Unger, Wilson Sawyer and Gerry White of Nortel Networks, Inc. Data Over Cable Division wrote the original versions of Section 3 and Section 5 through Section 10, as well as Appendices A through E, G, H, I, and K. Wilson Sawyer authored the original version of the Theory of Operations section of Chapter 8. John Ulm and Rick Pioso revised all of the state transition diagrams in Section 9.4 and authored the new Appendix L.

Steven Anderson, Tom Kolze, Victor Hou and Clive Holborow of General Instrument Corporation wrote Section 4 and Appendix F.

John Chapman of Cisco Systems wrote the Dynamic Service Change, Payload Header Suppression, and Dynamic Channel Change sections; extensively rewrote Appendix H, and contributed the idea for Unsolicited Grants with Activity Detection—including authoring Appendix M.

John Pickens and James Yee of Com21 made significant contributions to the Theory of Operations section of Chapter 8 and the QoS Parameters in Appendix C as well as updating Appendix E

Mike St. Johns of @Home Corporation spearheaded the addition of IGMP support to the specification. The contributions of Howard Abramson and Paul Gray of Motorola Corporation provided an additional compatible.

CableLabs and the cable industry as a whole are grateful to these individuals and organizations for their contributions.

This page intentionally left blank.

Appendix R Revisions

R.1 ECNs included in SP-RFIv1.1-I02-990731

Table R-1. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-99019	07/07/99	James Yee
rfi-n-99025	05/19/99	Rusty Cashman
rfi-n-99035	06/02/99	Rusty Cashman
rfi-n-99039	06/23/99	Rusty Cashman
rfi-n-99040	06/23/99	Rusty Cashman
rfi-n-99043	06/30/99	Doug Jones
rfi-n-99048	07/07/99	Clive Holborow
rfi-n-99049	07/07/99	Rusty Cashman
rfi-n-99050	07/07/99	Doug Jones
rfi-n-99051	07/07/99	Doug Jones
rfi-n-99052	07/07/99	Doug Jones
rfi-n-99053	07/07/99	Doug Jones
rfi-n-99054	07/07/99	Doug Jones
rfi-n-99056	07/07/99	Doug Jones

R.2 ECNs included in SP-RFIv1.1-I03-991105

Table R-2. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-99058	8/11/99	Bill Kostka
rfi-n-99062	8/11/99	Rusty Cashman
rfi-n-99064	8/4/99	James Yee
rfi-n-99067	8/25/99	Diego Mazzola / John Pickens
rfi-n-99070	9/22/99	Mike Patrick
rfi-n-99071	9/8/99	Clive Holborow
rfi-n-99072	9/15/99	Clive Holborow
rfi-n-99073	9/29/99	Jeff Lee
rfi-n-99075	9/22/99	Tim Dailey
rfi-n-99078	9/29/99	Jonathan Min
rfi-n-99080	9/29/99	Doug Jones
rfi-n-99081	9/29/99	Lisa Ruby
rfi-n-99084	10/13/99	Tom Quigley
rfi-n-99085	10/24/99	John Ulm
rfi-n-99086	10/20/99	Rusty Cashman
rfi-n-99087	10/20/99	Andrew Sundelin
rfi-n-99089	11/3/99	Clive Holborow

R.3 ECNs included in SP-RFIv1.1-I04-000407

Table R-3. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-99092	11/24/99	Yoav Hebron
rfi-n-99095	12/1/99	Tim Dailey
rfi-n-99097	12/1/99	Kaz Ozawa
rfi-n-99098	12/1/99	Pak Siripunkaw
rfi-n-99101	12/15/99	Victor Hou
rfi-n-99102	12/15/99	Victor Hou
rfi-n-99103	2/16/00	Lucy Pollak
rfi-n-99104	2/16/00	Rusty Cashman
rfi-n-00006	3/15/00	Tim Dailey
rfi-n-00008	3/29/00	Pak Siripunkaw
rfi-n-00009	3/1/00	Alan Wellington
rfi-n-00010	3/1/00	Wilson Sawyer / Julie McGray
rfi-n-00014	3/15/00	Lisa Ruby
rfi-n-00015	3/15/00	Julie McGray / Lisa Ruby
rfi-n-00017	3/15/00	Clive Holborow
rfi-n-00018	3/15/00	Tim Dailey
rfi-n-00019	3/15/00	Mike Patrick
rfi-n-00020	3/15/00	James Fletcher
rfi-n-00022	3/15/00	Greg White
rfi-n-00023	3/15/00	Greg White

R.4 ECNs included in SP-RFIv1.1-I05-000714

Table R-4. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-00021	4/12/00	John Chapman
rfi-n-00029	4/05/00	Andre Lejeune
rfi-n-00032	5/03/00	Michele Damoiseau
rfi-n-00033	5/03/00	Luc Martens
rfi-n-00036	5/31/00	Tim Dailey
rfi-n-00043	5/31/00	Harold Roberts
rfi-n-00046	6/07/00	Clive Holborow
rfi-n-00047	6/14/00	Clive Holborow
rfi-n-00050	6/21/00	Clive Holborow / Cliff Epps
rfi-n-00051	6/21/00	Rusty Cashman
rfi-n-00052	6/21/00	Chris Leano
rfi-n-00053	6/21/00	Richard Catchpole
rfi-n-00055	6/21/00	Greg White
rfi-n-00056	6/21/00	Doug Jones
rfi-n-00057	6/21/00	Doug Jones
rfi-n-00058	6/21/00	Doug Jones
rfi-n-00059	6/21/00	Greg White

R.5 ECNs included in SP-RFIv1.1-I06-001215

Table R-5. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-00070	8/23/00	Rusty Cashman
rfi-n-00072	8/30/00	Greg White
rfi-n-00073	9/27/00	Greg White
rfi-n-00075	9/27/00	Kartik Chandran
rfi-n-00079	10/4/00	Andrew Sundelin
rfi-n-00082	9/27/00	Margo Dolas
rfi-n-00086	10/11/00	Thomas Shurtliff
rfi-n-00088	10/25/00	Lisa Ruby
rfi-n-00089	10/25/00	Michael Patrick
rfi-n-00095	11/1/00	Rusty Cashman
rfi-n-00099	11/1/00	Clive Holborow
rfi-n-00105	11/15/00	Greg Nakanishi
rfi-n-00112	11/22/00	Greg White
rfi-n-00113	11/22/00	Greg White
rfi-n-00114	11/22/00	Greg White
rfi-n-00115	11/22/00	Greg White

R.6 ECNs included in SP-RFIv1.1-I07-010829

Table R-6. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-00120	12/06/00	Doug Jones
rfi-n-00121	12/06/00	Junming Gao
rfi-n-01012	03/14/01	Greg White
rfi-n-01027	04/04/01	Rusty Cashman
rfi-n-01031	04/04/01	Yoav Hebron Victor Hou
rfi-n-01039	04/18/01	Greg White
rfi-n-01059	05/23/01	Greg White
rfi-n-01060	06/13/01	Margo Dolas
rfi-n-01071	06/20/01	Greg White

R.7 ECNs included in SP-RFIv1.1-I08-020301

Table R-7. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-01042	11/07/01	Jeremy Gasper
rfi-n-01049	08/08/01	Dmitrii Loukianov
rfi-n-01051	11/07/01	Clive Holborow
rfi-n-01074	08/15/01	Clifton Epps
rfi-n-01085	10/31/01	Lucy Pollak
rfi-n-01086	11/07/01	Aryeh Elkon
rfi-n-01092	11/28/01	Chad Griffiths Nabil Valji
rfi-n-01098	12/19/01	Efrat Zeharhary Margo Dolas
rfi-n-01100	11/28/01	André Lejeune
rfi-n-01102	12/05/01	Margo Dolas
rfi-n-01105	12/19/01	Margo Dolas
rfi-n-02003	02/13/02	Elias Nemer
rfi-n-02005	02/13/02	Kaz Ozawa
rfi-n-02010	02/13/02	Sumanth Channabasappa/ Sujan Kumar

R.8 ECNs included in SP-RFIv1.1-I09-020830

Table R-8. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-02007	02/20/02	Margo Dolas
rfi-n-02013	02/20/02	Eduardo Cardona
rfi-n-02019	03/06/02	Kazuyoshi Ozawa
rfi-n-02030	04/03/02	Greg White
rfi-n-02041	03/20/02	Greg White
rfi-n-02043	03/27/02	Yael Barzilai
rfi-n-02051	05/02/02	Mike Grimwood
rfi-n-02073	05/01/02	Jason Tessier
rfi-n-02074	05/08/02	Efrat Zeharhary
rfi-n-02086	06/12/02	André Lejeune
rfi-n-02091	06/12/02	Rusty Cashman
rfi-n-02092	7/31/02	Zvi Shteingart
rfi-n-02109	06/26/02	Victor Hou
rfi-n-02115	06/19/02	Matt Schmitt
rfi-n-02143	07/07/02	Greg Nakanishi
rfi-n-02144	08/14/02	Greg White
rfi-n-02169	08/14/02	Greg White

R.9 ECNs included in SP-RFIv1.1-I10-030730

Table R-9. Incorporated ECN Table

ECN	Date Accepted	Author
rfi-n-02162	08/21/02	Andre Lejeune
rfi-n-02163	08/21/02	Gordon Li
rfi-n-02164	09/04/02	Andre Lejeune
rfi-n-02177	10/09/02	Victor T. Hou
rfi-n-02186	11/06/02	Joe Andonieh
rfi-n-02209	11/20/02	Greg White
rfi-n-02211	12/11/02	John Pickens
rfi-n-02225	12/11/02	David Pullen
rfi-n-02240	02/05/03	Rusty Cashman
rfi-n-03007	02/05/03	Efrat Zeharhary
rfi-n-03049	05/14/03	Lucy Pollak
rfi-n-03062	07/02/03	Greg White
rfi-n-03072	07/02/03	Shane Keck
rfi-n-03074	07/02/03	Lina Nakhle
rfi-n-03076	07/02/03	Greg White

R.10 ECNs included in SP-RFIv1.1-C01-050907

Table R-10. Incorporated ECN Table

ECN	Date Accepted	Author
RFIv1.1-N-05.0242-1	8/31/05	Greg White